MIT Technology Review

VOL. 121 NO. 2 MARCH/APRIL 2018 US $6.99/CAN $7.99

3-D metal printing

Babel-fish earbuds

The sensing city

AI for everyone

Dueling neural networks

Materials' quantum leap

Zero-carbon natural gas

Perfect online privacy

Artificial embryos

and

Genetic fortune-telling

10 Breakthrough Technologies 2018

BLOCKCHAIN

YOU'VE HEARD THE TERM, NOW UNDERSTAND THE IMPACT.

Meet the Thought-leaders and Innovators

Brian Behlendorf
Hyperledger Project
Blockchain: State of the State

Neha Narula
MIT Digital Currency Initiative
Examining Open, Trustless Digital Payments and Contracts

Rod Garratt
University of California at Santa Barbara
Cryptocurrencies and Global Monetary Policy

Gary Gensler
Commodity Futures Trading Commission
Ethics and Governance in the Blockchain Era

Frank Yiannas
Walmart
Industry Impact: Food Safety in the Global Supply Chain

BUSINESS OF BLOCKCHAIN

MIT Technology Review

mit media lab

APRIL 23, 2018 - CAMBRIDGE, MA

Join the conversation at Technologyreview.com/blockchain2018

From the Editor

Gideon Lichfield is editor in chief of **MIT Technology Review.**

Our annual list of 10 world-changing technologies invariably defies attempts to find an overarching theme. But a look back at the past few years shows a trend: we're including more and more advances in artificial intelligence.

We've featured surprise modeling, a form of machine learning (2008); Siri (2009); deep learning (2013); neuromorphic chips (2014); conversational interfaces (2016); robots that teach each other (2016); self-driving trucks (2017); and reinforcement learning (2017).

Algorithms that learn have been around for decades, so why this sudden flowering? It's thanks in part to better algorithms, but mostly to an explosion in the quantity of data available for training them—from photos to disease statistics to online shopping patterns—and to new kinds of chips that can better handle their massive processing needs.

This year's list again contains two AI entries. Generative adversarial networks, or GANs (page 42), are AIs pitted against one another in an evolutionary arms race, which speeds up the pace of learning, sometimes by orders of magnitude. And cloud-based AI (page 40) makes deep-learning algorithms as ubiquitous and accessible as blogging software.

Combined, these two innovations could put far more AI power in the hands of far more people. As is usually the case with a powerful technology, this cuts both ways. It promises to turbocharge scientific research and economic productivity. But it may also allow almost anyone to craft convincing fake images and videos that further erode society's ability to distinguish truth from lies. Ian Goodfellow, the inventor of GANs, is the rare technologist who is actively working on countering possible abuses of his invention, as Martin Giles relates in a profile of him (page 48).

There's a similar good-or-evil dichotomy in some of the other technologies on this year's list. As Antonio Regalado reports (page 54), genetics is evolving from the search for "the gene for X" to statistical analyses of big genomic databases that can now reveal "the thousand genes for X." That will help identify people at high risk of things like heart disease and Alzheimer's, but it will also predict traits such as height and IQ. What will we do with such knowledge? Synthetic human embryos (page 39) will help medical researchers study the earliest stages of life, but when does it become unethical to grow one? Alphabet's Sidewalk Labs plans to turn a Toronto neighborhood into a smart-city experiment (page 60), stuffing it with sensors to scoop up data on its residents' every movement. Will it, asks Elizabeth Woyke, become a shining example of how to use big data for public good, or a privacy nightmare?

Luckily, perhaps, not everything on our list is so ethically fraught. A natural-gas power plant that doesn't emit carbon dioxide; a 3-D printer for metal; simultaneous-translation devices that nestle in your ear like the Babel fish of the cult classic *Hitchhiker's Guide to the Galaxy*. All these and more are in this, our 17th annual list of the 10 breakthrough technologies. I hope you enjoy it.

NEPHI NIVEN

AUTONOMOUS VEHICLE TECHNOLOGY

AUTONOMOUS VEHICLES: ARE YOU READY FOR THE NEW RIDE?

Automakers are joining with Google, Uber, and high-profile start-ups to harness the technologies that will power next-generation autonomous vehicles.

THE SELF-DRIVING CAR revolution is about to shift into overdrive. Just look in the back lots of South Boston's tech corridor, on the streets of Pittsburgh, in the prefab test facilities at the University of Michigan's Mcity, and throughout a smattering of open highways.

Car companies are joining with tech giants like Google, Uber, and prominent start-ups to develop next-generation autonomous vehicles that will alter our roads and throughways and lay the groundwork for future smart cities. They're harnessing technological advances such as machine learning, Internet of Things (IoT), and the cloud to accelerate development.

More significantly, autonomous vehicles will advance the industry disruption set in motion by popular ride-sharing services like Uber and Lyft. "Autonomous vehicles will help bring the city back to what it was—for people," says Ryan Chin, co-founder of Optimus Ride, a Boston-based start-up working on self-driving technologies. As autonomous vehicles gain traction, Chin envisions an opportunity to remake the city landscape by consolidating parking, reclaiming land for parks, reducing urban congestion and traffic, and promoting overall highway safety.

Eventually, all self-driving cars will employ some combination of sensors, cameras, radar, high-performance GPS, Light Detection and Ranging (LIDAR), artificial intelligence (AI), and machine learning to achieve their respective levels of autonomy.

Societal Implications

While the era of the connected car is still in its infancy, there is widespread optimism that it's the wave of the future. By 2040, autonomous vehicles are expected to comprise around 25 percent of the global market.

The rise of the connected vehicle has far-reaching societal implications, from environmental benefits to improved safety.

As more autonomous vehicles hit the streets, city planners will accelerate plans to modernize highways and thoroughfares with smart technology for road signs, traffic lights, and merge lanes—all in an effort to reduce congestion and increase public safety. "For decades, we thought we'd get to smart highways by making highways smart, but there's been a huge chicken-and-egg problem," says David Tennenhouse, chief research officer for VMware. "With autonomous vehicles that can navigate existing infrastructure and talk to each other as they do [through vehicle-to-vehicle communications], we can get more efficiency out of roads."

Technology Enablers

For both the self-driving cars and the smart roadway systems, endpoint telemetry, smart software, and cloud are essential enablers—and cloud-based networking and connectivity are an important part of the mix. Advanced algorithms, AI, and deep learning systems are central to ensuring that self-driving cars can quickly and automatically adapt to changing scenarios.

Beyond the specific components, scalability of cloud computing infrastructure along with intelligent data management and transmission capabilities are indispensable for ensuring all of the right information is processed properly and securely.

While this spate of advanced technologies is providing an on-ramp to autonomous vehicle development, Chin, Tennenhouse, and other experts admit it will be a while before the self-driving revolution forever alters the transportation landscape. "The final frontier for autonomy is urban driving under extreme weather conditions, and that's going to take some time," says Chin.

For more information on technology that enables connected transportation, visit **vmware.com/go/connectedcar**.

vmware®

Contents

March/April 2018

Front

VIEW

UPFRONT

ON THE COVER

Illustrations by Derek Brahney

10 BREAKTHROUGH TECHNOLOGIES

p. 36

Q&As

Back

REVIEWS

DEMO

BACK PAGE

AUTONOMOUS VEHICLE TECHNOLOGY

DRIVERLESS VEHICLES WILL TALK TO US

Envisioning a not-so-distant future where our cars can detect our locations, our to-do lists, and our heart rates.

WHEN WE SWITCHED from landlines to smartphones about 10 years ago, there was a fundamental shift in how we understood communication technology. Until that point, we'd been using phones pretty much the same way for over 100 years.

A similar shift is about to happen with cars. With autonomous vehicles on the horizon, we can expect our relationship with cars—and our cars' relationship to the world—to change, and soon.

Automated Vehicles Will Talk to Roads and Traffic Signals

Self-driving cars will be capable of seamless parallel parking. But that's not the only thing that will impact parking innovation, according to Johanna Zmud, the Director of the National Office at the Texas A&M Transportation Institute.

"If I'm an urban dweller and I have a self-driving car, the car may not have to be parked in an expensive urban parking lot," she says. Instead, it could be stored in a more remote, less-bank-busting lot or suburban street, then be beckoned back using mobility capabilities.

Moreover, vehicle-to-infrastructure communication could enable driverless cars to communicate with traffic signals, says Zmud. Not only would it allow cars to follow travel rules, it could also optimize how cars move along roadways.

Automated Vehicles Will Talk to Us

Our interactions with driverless cars will look a whole lot like the ways in which we currently interact with computers, phones, and the internet—it will just go a step further. This kind of interaction is called a digital/physical hybrid network. Integrated Roadways CEO and CTO Tim Sylvester offers an example: It's similar to the ad suggestions we get from Google and Facebook, but with immediate, physical solutions that allow us to go places, make purchases, and interact with other people.

Automated Vehicles Will Talk to Everything

When cars are connected, they become vulnerable to security breaches. The risk isn't any higher than other hacking risks, says Zmud—your computer is just as likely to be compromised, if not more so. Nevertheless, the intersection between connectivity and safety presents a unique opportunity for engineers to get creative with their cybersecurity solutions.

Automakers could even develop technology that triggers automatic and total shutdowns of vehicles, says Sylvester. Or push software updates to cars without requiring the customer to take it to the dealer.

"Let's say several cars are reporting issues with their braking systems," says Mimi Spier, vice president of VMware's Internet of Things business. "The manufacturer could update the car's software to fix the problem in time to prevent accidents." The in-vehicle systems could then inform the driver that the software update successfully addressed specific issues.

Automated Vehicles Will Change Our Expectations

The changes that will come along with autonomous vehicles are a lot to consider—cars that talk to houses? Cars that suggest we buy flowers for our spouses? None of that fits in the paradigm of driving that we've been familiar with for over 100 years. Nevertheless, we've changed our expectations before, and we can do it again.

Learn more about the connected car

Download the five-part VMware "Connected Car Business Brief Series" for more information about the infrastructure, applications, and data analysis tools required for the autonomous and connected car of the future. Visit **vmware.com/go/car**.

vmware®

CUSTOMER SERVICE AND SUBSCRIPTION INQUIRIES

National
800-877-5230

International
903-636-1115

E-mail
customer_service@mittechnologyreview.info

Web
www.technologyreview.com/customerservice

MIT Records (alums only)
617-253-8270

Reprints
techreview@wrightsmedia.com
877-652-5295

Licensing and permissions
licensing@technologyreview.com

Technology Review
One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to equip its audiences with the intelligence to understand and contribute to a world shaped by technology.

AUTONOMOUS VEHICLE TECHNOLOGY

MIND-BLOWING TO MUNDANE: HOW TECH IS RESHAPING OUR EXPECTATIONS

Tech innovation is reshaping our expectations and transforming every industry. Simply put, tech is leaving the nest.

BY PAT GELSINGER

CONSIDER YOUR FIRST RIDE in a self-driving car. In the first 60 seconds, you may feel euphoria and deep fear: "Nobody's driving this thing!" In minutes two and three, there's fierce questioning: "What if someone swerves into my lane?" By the fourth minute, you're bored and checking your smartphone. The experience has gone from mind-blowing to mundane in four minutes flat.

In this time of seismic change, what's the most profound change of all? I believe it's our basic frame of reference—our everyday expectations.

Tech Is Leaving the Nest

Tech innovation is reshaping our expectations and transforming every industry you can name, from banking to media, healthcare to retail, and beyond. Practically every element of every sector is deeply and profoundly changing through these digital transformations.

Nowhere is this dynamic more evident than in the technology we use to care for premature babies. In the 1980s, babies born any earlier than 28 weeks of gestation were routinely offered only palliative care because their survival chances were so low. Today a baby born at 28 weeks has a 90 percent chance of survival, while babies born at 24 weeks have a 50/50 chance.

This summer I visited Children's Healthcare of Atlanta (CHOA) and their neonatal intensive care unit, a state-of-the-art facility. There, preterm babies depend on sophisticated software technology to breathe, until their heart and lungs can do the job on their own. Twenty years ago, not many predicted that IT software would be so critical to critical care.

Innovating in Cloud and Cybersecurity

In this era when tech innovation is transforming every industry, we in the tech industry have to recognize that we still have a lot of work to do. In particular, it's imperative that we innovate to improve two core building blocks of digital transformation—namely, the cloud and cybersecurity. Every CIO I talk to is looking for a better and simpler way to manage the critical apps and data that they're already running across multiple clouds.

We face a similar challenge in the realm of cybersecurity. As an industry, our goal is simple: We need to take all those disparate security point solutions and architect them into a secure infrastructure—one capable of tightly integrating all of the major security software components into a cohesive whole.

Over the past decade, cloud has changed the rules of the game and I expect the cloud model will remain dominant for the next five to ten years. With the Internet of Things taking root, we're going to see massive demand for compute power that extends out to the edge.

One Final Thought

It may feel like the pace of technology disruption and change these days is so dizzying that it could not possibly get any more intense. Yet here's the science fact: the pace of change right now is the absolute slowest it will be for the rest of your life. Fasten your seatbelts. It's going to be a fascinating ride.

For more on perspectives on technology innovation, visit **vmware.com/radius**.

vmware®

Views

ARTIFICIAL INTELLIGENCE

Tech companies should stop pretending AI won't destroy jobs

No matter what anyone tells you, we're not ready for the massive societal upheavals on the way.

I took an Uber to an artificial-intelligence conference at MIT one recent morning, and the driver asked me how long it would take for autonomous vehicles to take away his job. I told him it would happen in about 15 to 20 years. He breathed a sigh of relief. "Well, I'll be retired by then," he said.

Good thing we weren't in China. If a driver there had asked, I would have had to tell him he'd lose his job in about 10 years—maybe 15 if he was lucky.

That might sound surprising, given that the US is, and has been, in the lead in AI research. But China is catching up—if it hasn't already—and that rivalry, with one nation playing off the other, guarantees that AI is coming.

China will have at least a 50/50 chance of winning the race, and there are several reasons for that.

First, China has a huge army of young people coming into AI. Over the past decade, the number of AI publications by Chinese authors has doubled. Young AI engineers from Face++, a Chinese face-recognition startup, recently won first place in three computer-vision challenges—ahead of teams from Google, Microsoft, Facebook, and Carnegie Mellon University.

Second, China has more data than the US—way more. Data is what makes AI go. A very good scientist with a ton of data will beat a super scientist with a modest amount of data. China has the most mobile phones and internet users in the world—triple the number in the United States. But the gap is even bigger than that because of the way people in China use their devices. People there carry no cash. They pay all their utility bills with their phones. They can do all their shopping on their phones. You get off work and open an app to order food. By the time you reach home, the food is right there, hot off the electric motorbike. In China, shared bicycles generate 30 terabytes of sensor data in their 50 million paid rides per day—that's roughly 300 times the data being generated in the US.

Third, Chinese AI companies have passed the copycat phase. Fifteen years ago almost every decent startup in China was simply copying the functionality, look, and feel of products offered in the US. But all that copying taught eager Chinese entrepreneurs how to become good product managers, and now they're on to the next stage: exceeding their overseas counterparts. Even today, Weibo is better than Twitter. WeChat delivers a way better experience than Facebook Messenger.

And fourth, government policies are accelerating AI in China. The Chinese government's stated plan is to catch up with the US on AI technology and applications by 2020 and to become a global AI innovation hub by 2030. In a speech in October, President Xi Jinping encouraged further integration of the internet, big data, and artificial intelligence with the real-world economy. And in case you're wondering, these things tend not to be all talk in China—as demonstrated with its past policies promoting high-speed rail and the mass entrepreneurship and innovation movement. In comparison, things get bogged down in the US. Consider the way President Barack Obama's loan guarantee to solar-panel

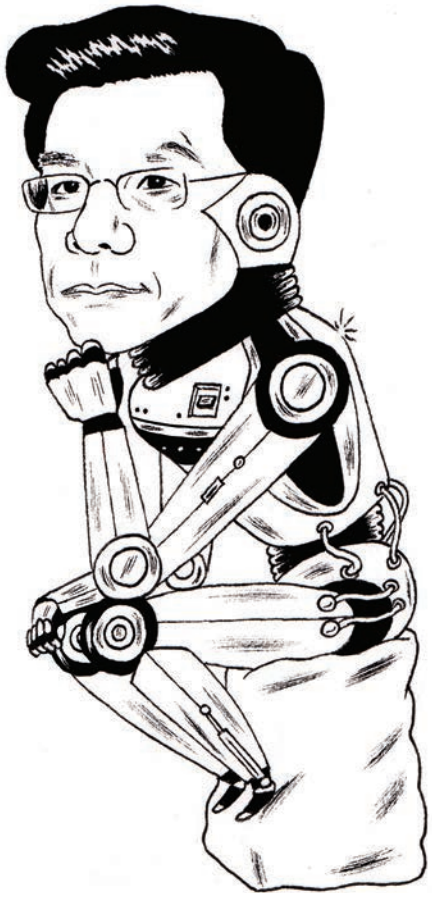

Kai-Fu Lee

ANDY FRIEDMAN

maker Solyndra was hammered as crony capitalism. Truckers are now appealing to President Donald Trump and Congress to stop testing of autonomous trucks.

The rise of China as an AI superpower isn't a big deal just for China. The competition between the US and China has sparked intense advances in AI that will be impossible to stop anywhere. The change will be massive, and not all of it good. Inequality will widen. As my Uber driver in Cambridge has already intuited, AI will displace a large number of jobs, which will cause social discontent. Consider the progress of Google DeepMind's AlphaGo software, which beat the best human players of the board game Go in early 2016. It was subsequently bested by AlphaGo Zero, introduced in 2017, which learned by playing games against itself and within 40 days was superior to all the earlier versions. Now imagine those improvements transferring to areas like customer service, telemarketing, assembly lines, reception desks, truck driving, and other routine blue-collar and white-collar work. It will soon be obvious that half of our job tasks can be done better at almost no cost by AI and robots. This will be the fastest transition humankind has experienced, and we're not ready for it.

Not everyone agrees with my view. Some people argue that it will take longer than we think before jobs disappear, since many jobs will be only partially replaced, and companies will try to redeploy those displaced internally. But even if true, that won't stop the inevitable. Others remind us that every technology revolution has created new jobs as it displaced old ones. But it's dangerous to assume this will be the case again.

Then there are the symbiotic optimists, who think that AI combined with humans should be better than either one alone. This will be true for certain professions—doctors, lawyers—but most jobs won't fall in that category. Instead they are routine, single-domain jobs where AI excels over the human by a large margin.

Others think we'll be saved by a universal basic income. "Take the extra money made by AI and distribute it to the people who lost their jobs," they say. "This additional income will help people find their new path, and replace other types of social welfare." But UBI doesn't address people's loss of dignity or meet their need to feel useful. It's just a convenient way for a beneficiary of the AI revolution to sit back and do nothing.

It will soon be obvious that half our tasks can be done better at almost no cost by AI. This will be the fastest transition humankind has experienced, and we're not ready for it.

And finally, there are those who deny that AI has any downside at all—which is the position taken by many of the largest AI companies. It's unfortunate that AI experts aren't trying to solve the problem. What's worse, and unbelievably selfish, is that they actually refuse to acknowledge the problem exists in the first place.

These changes are coming, and we need to tell the truth and the whole truth. We need to find the jobs that AI can't do and train people to do them. We need to reinvent education. These will be the best of times and the worst of times. If we act rationally and quickly, we can bask in what's best rather than wallow in what's worst.

Kai-Fu Lee is the founder and CEO of Sinovation Ventures and president of the Sinovation Ventures Artificial Intelligence Institute.


I'mpossible:
is ensuring your startup has the right backup
You act differently when you've got backup.
Like when your best friend is a
10th degree black belt. Or a tank.
You have a bit more spring in your step.
You go for it. Because what's the worst
that could happen? Well, in business,
it's lawsuits, cyber threats and identity theft.
Nothing you can't handle.
So go ahead, be your most
courageous self.
We've got your back.
Tailored business insurance for the courageous.
hiscox.com | 877-505-2048
HISCOX
encourage courage

MIT TECHNOLOGY REVIEW
VOL. 121 | NO.2

Upfront

Artificial intelligence that doubts itself

AI will make better decisions by embracing uncertainty.

SAIMAN CHOW

Upfront

The most powerful approach in AI, deep learning, is gaining a new capability: a sense of uncertainty. Researchers at Uber and Google are working on modifications to the two most popular deep-learning frameworks that will enable them to handle probability. This will provide a way for the smartest AI programs to measure their confidence in a prediction or a decision—essentially, to know when they should doubt themselves.

Deep learning, which involves feeding example data to a large and powerful neural network, has been an enormous success over the past few years, enabling machines to recognize objects in images or transcribe speech almost perfectly. But it requires lots of training data and computing power, and it can be surprisingly brittle.

Counterintuitively, perhaps, a sense of self-doubt offers one fix. The new approach could be useful in critical scenarios involving self-driving cars and other autonomous machines.

"You would like a system that gives you a measure of how certain it is," says Dustin Tran, who is working on this problem at Google. "If a self-driving car doesn't know its level of uncertainty, it can make a fatal error."

The work reflects the realization that uncertainty is a key aspect of human reasoning and intelligence. Adding it to AI programs could make them smarter and less prone to blunders, says Zoubin Ghahramani, a prominent AI researcher who is a professor at the University of Cambridge and chief scientist at Uber.

This may prove vitally important as AI systems are used in ever more critical scenarios. "We want to have a rock-solid framework for deep learning, but make it easier for people to represent uncertainty," Ghahramani told me recently over coffee one morning during a major AI conference in Long Beach, California.

During the same AI conference, a group of researchers gathered at a nearby bar one afternoon to discuss Pyro, a new programming language released by Uber that merges deep learning with probabilistic programming.

The meetup in Long Beach was organized by Noah Goodman, a professor at Stanford who is also affiliated with Uber's AI Lab. With curly, unkempt hair and an unbuttoned shirt, he could easily be mistaken for a yoga teacher rather than an AI expert. Among those at the gathering was Tran, who has also contributed to the development of Pyro.

"You would like a system that gives you a measure of how certain it is," says Dustin Tran, who's working on the problem at Google. "If a self-driving car doesn't know its level of uncertainty, it can make a fatal error."

Goodman explains that giving deep learning the ability to handle probability can make it smarter in several ways. It could, for instance, help a program recognize things, with a reasonable degree of certainty, from just a few examples rather than many thousands. Offering a measure of certainty rather than a yes-or-no answer should also help with engineering complex systems.

And while a conventional deep-learning system learns only from the data it is fed, Pyro can also be used to build a system preprogrammed with knowledge. This could be useful in just about any scenario where machine learning might currently turn up.

"In cases where you have prior knowledge you want to build into the model, probabilistic programming is especially useful," Goodman says. "People will use Pyro for all sorts of things."

Edward, yet another programming language that embraces uncertainty, was developed at Columbia University with funding from DARPA. Both Pyro and Edward are still at early stages of development, but it isn't hard to see why Uber and Google are interested.

Uber uses machine learning in countless areas, from routing drivers to setting surge pricing, and of course in its self-driving cars. The company has invested heavily in AI, hiring a number of experts working on new ideas. Google has rebuilt its entire business around AI and deep learning of late.

David Blei, a professor of statistics and computer science at Columbia University and Tran's advisor, says combining deep learning and probabilistic programming is a promising idea that needs more work. "In principle, it's very powerful," he says. "But there are many, many technical challenges."

Still, as Goodman notes, Pyro and Edward are also significant for bringing together two competing schools in AI, one focused on neural networks and the other on probability.

In recent years, the neural-network school has been so dominant that other ideas have been all but left behind. To move forward, the field may need to embrace these alternatives.

"The interesting story here is that you don't have to think of these camps as separate," Goodman says. "They can come together—in fact, they are coming together—in the tools that we are now building."

You might even say they are getting smarter, in part, by learning what they don't know. —*Will Knight*

Slack hopes its AI will keep you from hating Slack

The fastest-growing business app is relying on machine-learning tricks to fend off a deluge of messages—as well as competition from Facebook and Microsoft.

If you work at one of the 50,000 companies that pay to use Slack for workplace collaboration, you probably spend hours on it, swapping information, bantering, and sharing files with your colleagues. It's a casual, flexible way to interact—you tap out brief messages in group chat rooms (called channels) instead of sending e-mail, and it feels more like a smartphone app than typical office software.

But while it can be an efficient way to collaborate, keeping up with Slack can become a full-time task, particularly when you return from a few days away and find thousands of status updates, scattered across dozens of channels. Slack estimates that the average user sends 70 messages per day. How can you know which are must-reads and which can be skipped?

Slack's solution: artificial intelligence. In early 2016, the startup hired Stanford-trained computer scientist Noah Weiss to make the platform smarter. Over the past year and a half, Weiss's group has used machine learning to enable faster, more accurate information searches within Slack and identify which unread messages are likely to matter most to each user.

Slack says its platform, which launched publicly in 2014, is the fastest-growing business application ever, with more than six million daily active users. The company also predicts it will be bigger in the workplace than e-mail by 2025.

But e-mail isn't its only competition. Facebook, Google, and Microsoft, with their large existing user bases, have all released office collaboration tools recently. Microsoft says that 125,000 organizations use Microsoft Teams, its group-chat platform, which is bundled free with some Office 365 plans. Facebook says that more than 30,000 organizations, including Walmart, use its Workplace by Facebook service.

These chat products deliver not only steady revenues from monthly and annual service fees, but also troves of data that show how people interact within companies and what types of files and applica-

Noah Weiss is head of Slack's AI team, the Search, Learning, and Intelligence group.

JACKSON KRULE

Upfront

tions they use to get work done. Slack's larger competitors also see an opportunity to increase usage of their existing software. Companies like Microsoft "will tie these tools in with their other enterprise-wide platforms," such as Office 365, says Jeffrey Treem, an expert on communication technologies at the University of Texas at Austin. "All of these large technology companies are pursuing this same space because it's a very rich market."

Slack is not worried. "We think we have a bunch of important advantages, among them traction in the market, sharp focus, and a really deep understanding of our users," says CEO and cofounder Stewart Butterfield.

To understand how Slack intends to improve work through AI, I visited the company's New York office, where the team is based. Weiss built the 19-person group by recruiting engineers, designers, and product managers from companies like Facebook, Google, and LinkedIn, many from big-data projects. He boasts a similar résumé: one of his first jobs, after studying computer science and economics at Stanford, involved developing display ads at Google. After three years, Weiss moved to the startup Foursquare, where he led the product analytics team.

At Slack, Weiss is applying what he learned at Google and Foursquare to refine search queries and give people recommendations when they open the app. The work incorporates multiple AI methods, including different types of machine learning and natural-language processing.

Some of the technology is already live. One feature shows which people within a company talk about particular topics most often in Slack and where those discussions take place. The information, which appears when users conduct searches in Slack, is meant to pinpoint subject experts so people can direct questions to their most knowledgeable and accessible colleagues. Another feature, added last year, evaluates all of a user's unread messages, across all Slack channels; highlights up to 10 of the ones its algorithms deem most important; and presents them in a single list.

Both innovations rely on a data structure that Weiss calls the "work graph." It essentially looks at companies that use Slack and analyzes how the people within them are interrelated, where in the app their discussions are taking place, and what topics are being discussed. If the term sounds familiar, it's because Google and Facebook have similar structures—the "knowledge graph" and the "social graph," respectively. But while Google studies public data and Facebook promotes the idea of a single, global network of relationships, Slack thinks of the work graph as specific to each company—a representation of how work is structured within it.

Employees may balk if they think they'll be assessed on their Slack participation.

"We've spent a lot of time building models that understand what you care about and what content you interact with," says Jerry Talton, who helps lead the team's technical work. "In the future, we'll take that same understanding and apply it to content you don't know about that could make you better at your job."

Another Slack goal is to help management keep a better eye on its employees. One new initiative crunches data to give executives a bird's-eye view of how employees are interacting.

"You'd be able to see what your European set of offices are paying attention to versus your US set of offices, or what people who have long tenure at your company are paying attention to versus people who are really new," Weiss says.

But does mining employee communications invade their privacy? Weiss says his team hopes to assuage concerns by parsing activity only in public Slack channels (rather than the private ones where people can conduct confidential conversations). He also says Slack won't turn on the feature unless companies request it.

Still, employees may balk, particularly if they think they will get assessed on the basis of how active or popular they are on Slack. Adam Waytz, who researches social psychology and ethics at Northwestern University's Kellogg School of Management, thinks the feature sounds invasive. "Given the increasing public unease about employers' control over their employees' lives and what gets said at work, this product could result in backlash or paranoia," he says.

Slack also needs to gain trust for its existing AI features. "AI can be tremendously beneficial in matching the right people with the right information to do the right tasks, but it's not a perfect solution," says Treem. "If you were relying on algorithms to get you the most important messages and you find out a week later that you missed something ... you're going to lose confidence in Slack's ability to do what you need it to do."

Weiss says algorithm tweaks by the AI team last year made searches 50 percent more successful, and also made people 30 percent more likely to accept suggestions about new Slack channels to join. If all goes as planned, the intelligence layer the team is building on top of Slack will morph into a digital assistant that can make people more productive.

Butterfield, the Slack CEO, sees AI as a long game. "I think what we have right now is good," he says. "In a couple of years, it will be very good. In about five years, it will be excellent. And in 10 years it will be impossible to work without it."

—Elizabeth Woyke

GoBoldly.com

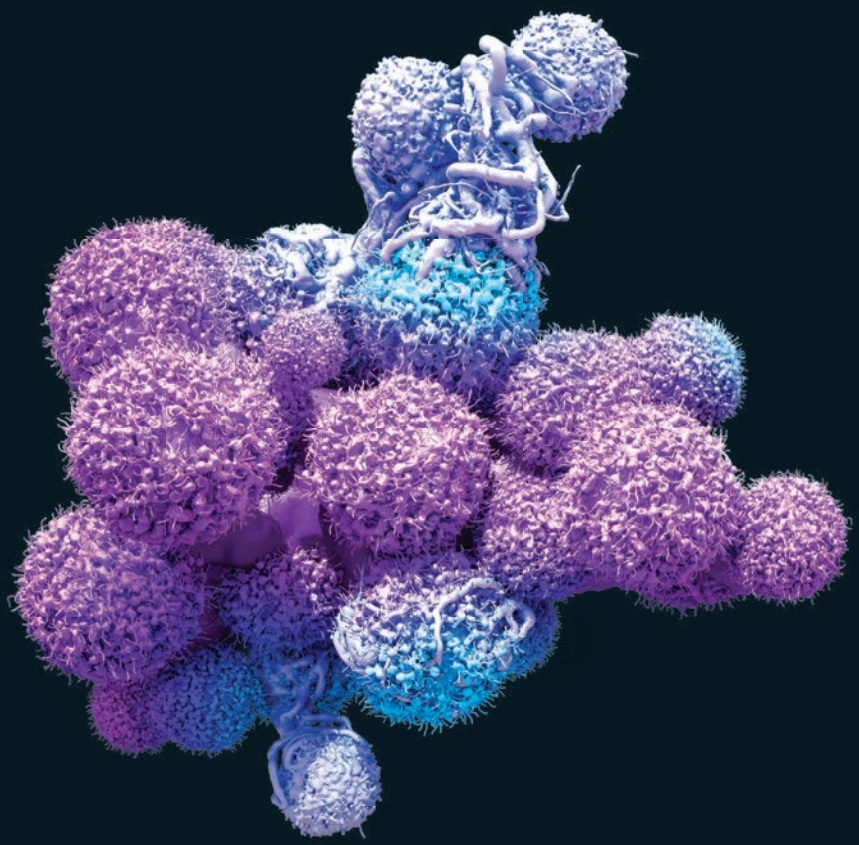

THERE WERE THOSE WHO BELIEVED THE BODY COULD NEVER FIGHT CANCER.

NEVER SAY NEVER.

Today, researchers are using immunotherapy treatments to stimulate the body's immune system to destroy invading cancer cells.
Welcome to the future of medicine. For all of us. GOBOLDLY™

America's Biopharmaceutical Companies℠

Upfront

How to get Wyoming wind to California

High-voltage direct-current transmission lines hold the key to slashing greenhouse gases.

Several miles south of Rawlins, Wyoming, on a cattle ranch east of the Continental Divide, construction crews have begun laying down roads and pads that could eventually underpin up to 1,000 wind turbines. Once complete, the Chokecherry and Sierra Madre project could generate around 12 million megawatt-hours of electricity annually, making it the nation's largest wind farm.

But how do you get that much wind power to where it's actually needed?

The Denver-based company behind the project hopes to erect a series of steel transmission towers that would stretch a high-voltage direct-current transmission (HVDC) line 730 miles across the American West. It could carry as much as 3,000 megawatts of Wyoming wind power to the electricity markets of California, Nevada, and Arizona. With the right deals in place, the transmission line could deliver solar-generated electricity back as well, balancing Wyoming's powerful late-afternoon winds with California's bright daytime sun.

The $3 billion TransWest Express Transmission Project is among a handful of proposed direct-current transmission lines in the United States, and one of the furthest along in the planning process. It underscores the huge promise of these high-capacity lines to unlock the full potential of renewable energy.

Transmission isn't sexy. It's basic infrastructure: long wires and tall towers. But a growing body of studies conclude that building out a nationwide network of DC transmission lines could help enable renewable sources to supplant the majority of US energy generation, offering perhaps the fastest, cheapest, and most efficient way of slashing greenhouse-gas emissions.

Developing these transmission lines, however, is incredibly time-consuming and expensive. The TransWest project was first proposed in 2005, but the developers will be lucky to secure their final permits and begin moving dirt at the end of next year.

There's no single agency in charge of overseeing or ushering along such projects, leaving companies to navigate a thicket of overlapping federal, state, county, and city jurisdictions—every one of which must sign off for a project to begin. As a result, few such transmission lines ever get built.

Direct current, in which electric charges constantly flow in a single direction, is an old technology. DC and AC—alternating current—were the subject of one of the world's first technology standards battles, pitting Thomas Edison against his former protégé Nikola Tesla in the "War of the Currents" starting in the 1880s.

AC won this early war, mainly because, thanks to the development of transformers, its voltage could be cranked up for long-distance transmission and stepped down for homes and businesses.

But a series of technological improvements have substantially increased the functionality of DC, opening up new ways of designing and interconnecting the electricity grid.

For the past two years, James McCalley, an engineering professor at Iowa State University, has been studying the best way to tie together those massive grid systems as part of the US Department of Energy's $220 million Grid Modernization Initiative.

One way to solve the problem is to expand existing "back-to-back" conversion stations to provide more east-to-west transmission capacity. These systems allow transmission between two grids, by converting the power to DC and then back to AC again at the point where they "cross the seam."

ADANI

Another approach adds three point-to-point transmission lines, running east to west, connecting the heart of each grid to that of the other. Yet another solution is a so-called "macro grid" of long DC transmission lines covering much of the country. It runs up the Florida panhandle, across the South, north to Seattle, east to Minneapolis, and back down to Louisiana, with several additional lines crisscrossing the West.

McCalley and his team developed models to simulate each of these scenarios over a 15-year period. They found that all three demonstrated a strong economic payoff, providing a benefit of at least $2.50 in savings for every $1 invested in the transmission system.

With direct-current lines, grid operators have more options for energy sources throughout the day, allowing them to tap into, say, cheap wind two states away during times of peak demand instead of turning to nearby but more expensive natural-gas plants for a few hours. The fact that regions can depend on energy from distant states for their peak demand also means they don't have to build as much high-cost generation locally.

A national direct-current grid could also help lower emissions to as much as 80 percent below 1990 levels within 15 years, all with commercially available technology and without increasing the costs of electricity, according to a study published earlier in *Nature Climate Change*.

The researchers produced an idealized transmission network that connected 32 nodes across the nation, linking hydroelectric power in the Pacific Northwest, solar in California, wind energy in the Southwest, and nuclear on the East Coast. Simply put, the system balances out the intermittency of renewable energy sources over long distances, meaning there's always reliable generation somewhere. Being able to tap into it from any corner of the nation lowers the cost of supplying energy at peak demand, reduces the amount of generation required in any single area, minimizes excess generation, and eliminates the need to develop expensive grid-scale storage systems.

"We're basically getting that big battery we want for free," says Christopher Clack, one of the lead authors of the study and chief executive of Vibrant Clean Energy.

There are already a handful of DC transmission lines in the US and a growing number of proposals, including the New England Clean Power Link, which would transport 1,000 megawatts of renewable power from Canada into New England. Houston's Clean Line Energy has at least a half-dozen proposals in various stages, including the Plains and Eastern Clean Line connecting western Oklahoma to markets in the Southeast, and the Grain Belt Express Clean Line stretching from Kansas to Indiana.

But all of these are moving through the approvals process at a dawdling pace. The TransWest developers, who have secured permission along the two-thirds of the line's path that lies on federal land since taking over the project in 2008, are still working to finalize approvals from states and private landowners.

Most developers and energy policy experts say what's needed to accelerate these projects is a federal authority with greater power to push them through. A report released by Stanford in October highlighted a number of possibilities, including granting the Federal Energy Regulatory Commission the same "siting authority" for transmission lines that it already has over natural-gas pipelines.

Dan Reicher, executive director of the Steyer-Taylor Center for Energy Policy and Finance at Stanford, who cowrote that report, says, "Without clear, predictable siting authority, it's going to be very difficult to build out an intelligent, comprehensive HVDC network." —*James Temple*

QUOTED

"A very dangerous and vulnerable cocktail."

—Security researcher Alexander Bolshev on mixing mobile apps and industrial control systems.

"We probably want to have some UX around that, because that's cool. It's very human."

—Jim Farley, Ford's executive vice president of global markets, on people's tendency to say "thank you" to driverless pizza delivery vehicles.

"We will eventually get to a point where we can detect cancer before it's otherwise visible."

—Len Lichtenfeld, deputy chief medical officer of the American Cancer Society, on the promise of blood tests that can catch cancer early.

BY THE NUMBERS

24

Number of hours the average American spends online each week, according to new data from USC Annenberg. That's up from 9.4 hours in the year 2000.

86

Number of people in China who have had their genes edited to treat their diseases, according to a report by the *Wall Street Journal.*

$90 billion

The amount carmakers have devoted to developing electric cars, according to Reuters. Some $19 billion of that was spent by US firms.

95 percent

Accuracy with which an AI assistant called Corti can recognize a heart attack during a 911 call, based on an early study in Denmark. Danish dispatchers can recognize a heart attack 73 percent of the time.

Upfront

23andMe launches a giant weight-loss study

The consumer DNA testing company seeks 100,000 volunteers to help it link genes to diet success.

Consumer DNA testing company 23andMe has begun what it terms a "massive study" into the genetic basis of weight loss. The company, based in Mountain View, California, has contacted 1.3 million of its customers with an offer to take part in the project by sticking to one of two diets or an exercise plan for three months, reporting back on whether their waistlines grew or shrank.

The crowdsourced study may prove to be the most comprehensive attempt yet to discern the links between people's genes and dieting success. 23andMe hopes what it learns will let it create predictive models that provide tailored weight-loss advice as part of its consumer genetic reports.

Already, consumers can pick from a dozen or more DNA tests that promise diet insights. But the tests have come under withering criticism from prominent doctors who say they're no better than the tips you'd get from a nutritionist or a friend at the gym. The advice might be okay; it's the DNA test that's a waste of money.

According to 23andMe, previous studies attempting to link DNA to dieting outcomes haven't had enough participants to zero in on genetic factors. Its new project will involve 10 to 50 times as many volunteers as previous work, says Geoffrey Benton, the company's head of health R&D.

The company holds DNA data on more than three million customers who have sent in saliva samples. That makes it one of the two or three largest biobanks in the world. After customers' DNA is analyzed, they receive reports about their geographic ancestry, how many Neanderthal genes they have, and a few hereditary health risks. Buyers also receive a prediction of their body mass based on their genes, a report telling them whether they have an inborn tendency to be heavier or thinner. The problem is 23andMe can't yet tell them what to do about it.

Starting last May, 23andMe began exploring whether it could convince its customers to carry out at-home experiments. It began with a pain tolerance test in which it asked people to see how long they could keep a hand in a bowl of ice water. That was followed by a sleep study in which about 6,000 volunteers were randomly assigned to change their behavior in specific ways, like avoiding coffee or agreeing not to look at a screen starting 30 minutes before bedtime. "We wanted to see if we could actually do an interventional trial from start to finish and do it remotely from 23andMe," Benton says.

With the new dieting study, 23andMe will randomly assign people to one of three plans. Some will avoid bread, cakes, and other carbohydrates. Another group will eat more fiber but shun animal fat. A third will eat as usual but add workouts to their week. They'll report back to the company about how often they have "cravings," whether they're stressed, and if they succeed in following the diets.

The company thinks people will have roughly the same results on all the plans. But it may be able to figure out whether there are genetic reasons why some people will lose 40 pounds and others gain 10 no matter which advice they follow.

With its diet study, 23andMe could demonstrate that its "platform" is suited to carrying out very large clinical trials of the type normally performed by research universities or drug companies. That could be commercially valuable, since 23andMe already sells genetic data to pharmaceutical companies and sometimes helps them locate people with specific diseases.

—Antonio Regalado

TO MARKET

Magic Leap One

Augmented-reality headset

COMPANY:
Magic Leap

PRICE:
Not yet available

AVAILABILITY:
This year

MAGIC LEAP

Augmented reality has advanced a lot in the past year. One of the most anticipated gadgets is from the secretive, massively funded (nearly $2 billion) Florida-based startup Magic Leap. It unveiled its first AR headset, Magic Leap One, at the end of 2017, saying it will ship in 2018. Styled sort of like a modern pair of steampunk goggles, the black headset is meant for developers, and the company says it will produce images that look great while being comfortable to wear—two goals that have eluded many AR headset makers thus far. Magic Leap has been working toward this release since 2011, but there are still a lot of details to be worked out, including price and release date. *—Rachel Metz*

BOUNDLESS CONNECTIVITY

GET PLUGGED IN TO OUR POWERFUL HIGH-TECH NETWORK.

From robotics and photonics to mobile technologies, and modeling and simulation, there is no doubt that Florida is a hotbed for high-tech excellence. That's why we're home to more than 30,000 tech companies that employ nearly 312,000 workers. We're also one of the largest exporters of high-tech goods. And considering we have over 2.5 million young, motivated millennials working here, we will be for a long time to come. Discover what a future in Florida means for your business at **floridathefutureishere.com/connectivity,** or call **877-YES-FLORIDA.**

FLORIDA | THE FUTURE IS HERE®

ENTERPRISE FLORIDA

Upfront

Behind South Korea's cryptocurrency boom

The country is a hub for trading virtual currencies, despite a government clampdown and North Korean cyberattacks.

South Korea is one of the world's most wired societies, with near-ubiquitous broadband access and blazing-fast internet speeds. Now the country is also becoming a hotbed for cryptocurrency trading. South Korea is the world's No. 3 market in Bitcoin trading, after Japan and the US, and the largest exchange market for Ether, Ethereum's cryptocurrency, accounting for more than 33 percent of its market share. The country is also home to two of the top 15 global digital-currency exchanges (Bithumb and Coinone), both of which have built walk-in centers where investors can conduct transactions in person. Overall, South Korea is believed to have about a million registered daily traders in virtual currency, equivalent to about one out of every 50 citizens.

The boom is worrying to the South Korean government, which is contemplating tighter regulations on cryptocurrency trading. Authorities are particularly concerned about a new method of fund-raising called an initial coin offering, in which companies create blockchain-based digital tokens that can be used to purchase a specific product or service in the future, and sell them publicly. Last September, the country's Financial Services Commission ordered a ban on ICOs. "Cryptocurrencies are neither money nor currency nor financial products," said the agency in a written statement at the time.

In South Korea, as in other countries, startups have been using ICOs to raise funds because the campaigns require little paperwork, let entrepreneurs solicit money directly from investors rather than rely on banks or venture capital firms, and enable founders to maintain total ownership of their companies. In September alone, South Korean startups raised about $89 million in sales of digital tokens, according to government data. When the FSC announced its ban in late September, 20 South Korean startups said they had planned to raise seed money through ICOs.

Retail investors aren't the only South Koreans excited about cryptocurrencies; some of the country's biggest corporations are pouring money into virtual-currency businesses. Nexon, one of South Korea's biggest video-game developers, is the leading shareholder in Korbit, the country's No. 3 cryptocurrency exchange. Dunamu, an affiliate of Kakao, a leading South Korean internet services company, recently launched a cryptocurrency exchange called Upbit. And the DB Group, another South Korean conglomerate, partnered with the local firm Sentbe in August to offer remittance payments in Bitcoin.

Even Samsung, South Korea's largest conglomerate, is getting involved. Last May, the company's IT solutions unit, Samsung SDS, announced a pilot project that will use blockchain technology to track shipping imports and exports, as well as the location of cargo shipments in real time.

South Korea's fervor for cryptocurrency is notable given that the country has an urgent reason to be skeptical: cyberattacks from North Korea. Hackers probably hailing from North Korea targeted officials at four South Korean Bitcoin exchanges last summer. The plots involved sending messages from stolen e-mail addresses and attaching malicious code that was identical to viruses previously proved to be of North Korean origin. Experts think the hackers were interested in Bitcoin because of its relative anonymity.

"The rampant use of digital currency offers both opportunities and risks," says Kim Kyung-soo, head of the Ethereum research center in South Korea. "Risk takers are attempting to make profits by delving into these high-volatility assets. But digital currencies could also be used as seed money to lift the next wave of technology developments." *—Yoochul Kim*

BRYAN FOUNTAIN

A contraceptive gel for men is about to go on trial

It's expected to deliver hormones more effectively than injections or pills.

After more than a decade of work, govern-ment researchers in the US are ready to test an unusual birth control method for men—a topical gel that could prevent the production of sperm.

And no, gentlemen, you don't rub it on your genitals.

The clinical trial, which begins in April and will run for about four years, will be the largest effort in the US to test a hormonal form of birth control for men.

Currently, the most effective options men have for birth control are condoms or a vasectomy. In the last major study of a hormonal male contraceptive, which took place in Europe from 2008 to 2012, participants received injections of hormones every two months. The shots suppressed sperm production and prevented the men's female partners from getting pregnant, but they also gave men severe mood swings and other serious side effects.

The new gel contains two synthetic hormones, testosterone and a form of progestin. Progestin blocks the testes from making enough testosterone for normal sperm production. The replacement testosterone is needed to counteract the hormone imbalances the progestin causes but won't make the body produce sperm.

More than 400 couples will participate in the study, which will take place at sites in the US, the UK, Italy, Sweden, Chile, and Kenya. Men in the trial will take home a pump bottle of the gel and rub about half a teaspoon of it on their upper arms and shoulders every day. The gel dries within a minute.

"It's not a lot of effort. It's just remembering to use it every day," says Diana Blithe, program director for contraception development at the National Institute of Child Health and Human Development, part of the National Institutes of Health.

The gel can suppress sperm levels for about 72 hours, so if men forget a dose, "there is a bit of forgiveness," says Régine Sitruk-Ware, distinguished scientist at the Population Council, a nonprofit for reproductive health that is sponsoring the trial alongside the NIH.

Men will use the gel for at least four months while their partners also use some form of female contraception. Researchers will monitor the men's sperm levels, which need to drop to less than one million per milliliter to effectively prevent pregnancy, according to Blithe. Once the sperm count is low enough, the women will go off their birth control. The couples will then use the contraceptive gel as their only form of daily birth control for a year.

The method was shown to be effective in an initial six-month study. But it involved two types of gels that had to be applied to different parts of the body, so Blithe's NIH team worked with researchers at the Population Council to reformulate the hormones and combine them into one gel.

Still, the question is: will men use it?

Historically, there hasn't been much interest from pharmaceutical companies in a male contraceptive. Running clinical trials takes years and is hugely expensive, so it's a risky endeavor when lots of options for female contraception already exist.

But researchers like Sitruk-Ware think views are changing, and that men, especially younger men, will be open to using a contraceptive drug. "This is about gender equity," she says. "Men would also like to be able to regulate their own fertility and not be forced into fatherhood."

Men's attitudes on their role in contraception vary by country, but a 2010 survey indicated that at least 25 percent of men worldwide would consider using a hormonal contraceptive. *—Emily Mullin*

TO MARKET

Luxturna

Gene therapy for blindness

COMPANY:
Spark Therapeutics

PRICE:
$425,000 per eye

AVAILABILITY:
Now

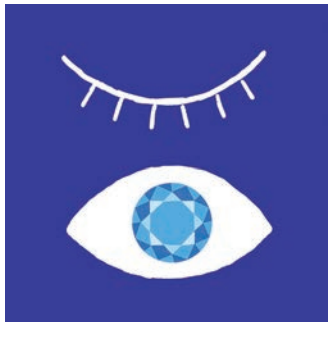

TK LEONARD GRECO

Spark Therapeutics recently announced the price tag for Lux-turna, the only FDA-approved gene therapy for an inherited type of blindness. That price is $425,000 per eye—making it the most expensive dose of a drug in the US. At $850,000 for both eyes, that's nearly nine times the median net worth of an American family. The therapy is meant to restore eyesight in people with a specific genetic mutation that causes degeneration of the retina. Spark hasn't said how many people that is, but it may be fewer than 30 patients a year. Is it worth it? While the therapy does stop eyesight from getting worse and even restores some sight, it's not a cure, and no one's sure how long the effects last.*—Emily Mullin*

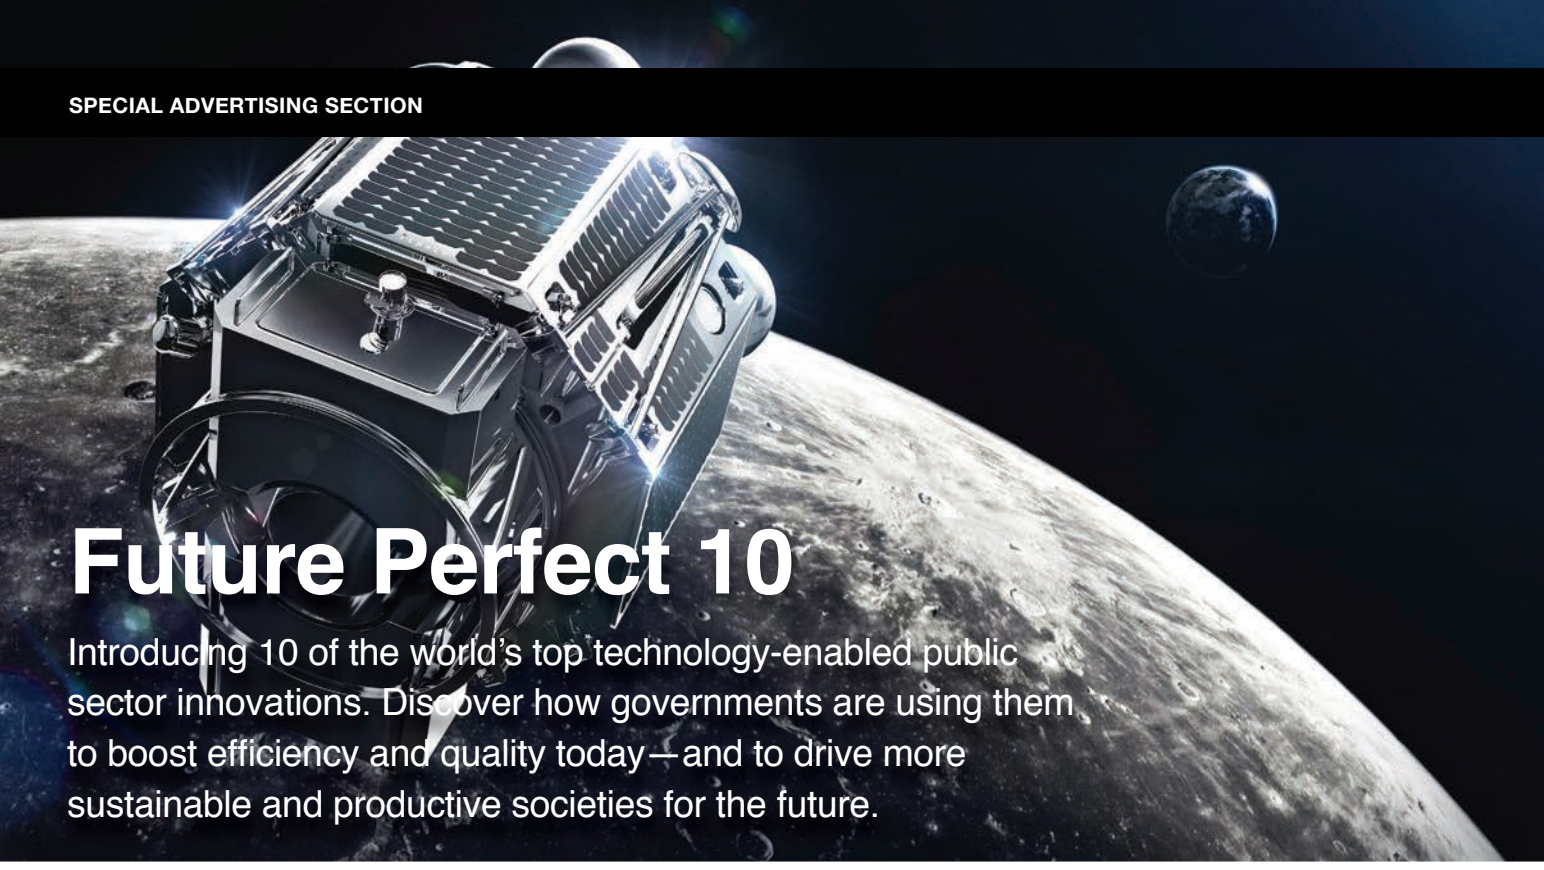

Future Perfect 10

Introducing 10 of the world's top technology-enabled public sector innovations. Discover how governments are using them to boost efficiency and quality today—and to drive more sustainable and productive societies for the future.

World governments have an insatiable need to make incremental improvements in the services they deliver while at the same time tackling the longer-term public sector challenges that limit their socioeconomic potential. As leading nations have discovered, the right use of digital technology is the pathway to achieving both objectives. Successful nations have technology strategies that are both aggressively practical and profoundly aspirational.

A breakthrough global public-private partnership recently collaborated to identify the top 10 government technology (govtech) initiatives worldwide that will contribute the most to sustainable improvements in 192 countries' economies and societies. MIT Technology Review Insights and the United Arab Emirates-based World Government Summit (WGS), a non-government organization promoting future-oriented public-private dialogue, worked together to select the most interesting govtech implementations from among the world's more innovative initiatives. The WGS reviewed the short list of 10 best-in-practice technology projects in order to select three recipients of its annual GovTech Prize—a recognition of their achievement in creating smart solutions which can inspire all digitally-minded governments.

The Government of Things

The growth in consumer use of digital devices creates a great opportunity for governments and societies. Using data generated from the Internet of Things (IoT), utility companies and transportation departments are improving the way they deliver services, for example:

1 SJ Railways—Sweden's state-owned national railway—is moving well beyond e-tickets by conducting the world's first NFC payment trial using microchips literally embedded in its customers, by offering the option to place a biometric chip under the skin in passengers' hands.

2 The Government of India's Aadhaar Program is the world's largest biometric identity card network, serving more than 1.2 billion citizens and disbursing nearly $40 billion in social welfare services annually.

Blockchain Party: Using Digital Trust to Bolster Civic Institutions

Blockchain applications—bitcoin is one—create systems of trust using digital ledgers, limiting regulatory oversight and licensing. In 2017, many governments got serious about using blockchain to enhance the authority and efficacy of their own public institutions, particularly those in emerging markets with traditionally weak governance, for instance:

3 Ghana's Ministry of Lands and Natural Resources is working with technology developer Bitland to create a complete

National governments often avoid digital transformation initiatives deemed too cutting edge. However, some forward-looking governments are using experimental new applications.

and permanent national land registry based on blockchain tokens.

4 The UAE's Smart Dubai Office is creating a blockchain-based paper-free digital transaction platform for its entire city government, to remove 100 million official documents from its processes. More than 20 blockchain use cases are now being pilot tested, in areas ranging from health and education to transportation and energy, aimed at making Dubai "the world's happiest and smartest city."

Govtech: Reaching for the Moon While Staying Grounded

National governments often avoid digital transformation initiatives deemed too cutting edge, as compliance issues, budget constraints, and the pressing demands of constituencies make far-ranging, blue-sky technology experiments impractical, if not politically risky. However, some forward-looking governments are using experimental new applications to address current service delivery challenges, or issues impacting economic sustainability.

5 The U.K.'s National Health Service is enlisting artificial intelligence (AI)- and analytic-enabled digital mental health providers such as Cambridgeshire-based Ieso Digital Health to increase access to cognitive therapy services to 1.5 million citizens by 2021. Ieso offers cognitive behavioral therapy via text-based messaging that takes place in a secure virtual room. The company collects data from sessions and uses machine learning to continue training its therapists in best practices.

6 Tanzania's Ministry of Agriculture, Food Security and Cooperatives distributes DNA sequencing devices to cassava farmers and researchers to aid in their efforts to eradicate whitefly-borne diseases. The farmers and scientists are using these portable, real-time devices to identify the strain of virus destroying their crops, allowing them to proactively fight the diseases.

7 The Development Bank of Japan, a government-owned bank, is the lead investor in ispace Inc., a commercial space launch company with ambitious plans to develop lunar infrastructure, financed in part by showing advertisements on the moon, as a first step to transforming Japan's economic development. The goal is to focus on the harmony between the earth and the moon, and on how the earth and the moon can become "one ecosystem."

Big Cities, Bigger Data

Collecting and parsing huge amounts of data is the lifeblood of the AI ecosystem. Two notable govtech initiatives, in Australia and China, are leveraging the wealth of data created by urban governments and municipal service providers, to not only make their cities more livable, but to make them R&D labs for their respective country's emerging AI industries.

8 Korea's Ministry of Land, Infrastructure and Transport has launched K-City, one of the world's largest autonomous-driving technology R&D facilities.

9 The Hangzhou Municipal Government has collaborated with tech giant Alibaba to create City Brain, a cloud-based AI platform to analyze traffic patterns, reducing congestion levels 10 percent this year by rerouting traffic around backups.

10 Australia's Minister for Cities and Digital Transformation launched the National Cities Performance Framework, an open data-based policy and planning tool to accelerate the country's Smart Cities rollout.

These leading govtech initiatives have helped deliver higher quality service to citizens and residents, but more importantly, they fundamentally change societal outcomes, allowing people to live safer, more productive, and more fulfilling lives.

ABOUT THE WORLD GOVERNMENT SUMMIT

The World Government Summit (WGS) is a UAE-based thought leadership organization which maintains a future-oriented dialogue between the private and public sectors of 140 countries on the best ways that governments can improve the lives of their citizens. This collaboration culminates each year in a global summit that serves as an incubator and an impetus shaping future global trends, with more than 150 distinguished thinkers, officials, and business leaders participating.

The top 10 initiatives that were selected comprise the short list for the WGS' Best Government Emerging Technologies Award, which acknowledges governments that experiment with emerging technologies to create greater societal value and revolutionize our existence. For more information, see www.worldgovernmentsummit.org; #WorldGovSummit

TONY LUONG

Q+A

Ethan Zuckerman

In the past, if you wanted to change the world, you had to pass a law or start a war. Now you create a hashtag. Ethan Zuckerman studies how people change the world, or attempt to, by using social media or other technological means. As director of the Center for Civic Media at MIT and an associate professor at the MIT Media Lab, he tries to help his students make sense of these issues. Zuckerman is also writing a book about civic engagement during a time when we have a lot less trust in institutions—government, businesses, banks, and so on. Maybe that lack of trust is reasonable. After all, we've spent the last decade-plus slowly turning our data over to large corporations like Facebook and Google without quite realizing we were doing it. Zuckerman knows what it's like to build technology that pisses a lot of people off. Back in the 1990s he created what became one of the most hated objects on the internet: the pop-up ad. The aim was to show an ad on a web page without making it look as though the advertiser necessarily endorsed the content on the page. "Our intentions," he later wrote in an apology to the internet at large, "were good." Zuckerman spoke with *MIT Technology Review* about how social media started controlling us rather than the other way around.

By Rachel Metz

RM **How are people using technology—rather than, say, lobbying for laws to be passed—to force change in new ways?**

EZ We used to make change mostly using law as our primary lever. Now we use the legal lever less; we use the levers of norms of markets and technology more often. #MeToo is an example of a norms-based campaign. It's basically saying, "We're going to challenge how people talk about sexual assault and sexual harassment." And once we change that norm, there's other legal pieces, market pieces, that'll come into play. But at its heart it's trying to change how we have certain conversations. The point in all of these is that if you can't get social change done through the traditional model of civics, there is a whole new set of tools, and people are starting to learn how to use these things.

RM **But social networks like Facebook and Twitter control, or at least direct, the information we see by using algorithms to filter what we see in our feed. You worked with two colleagues—Chelsea Barabas of MIT's Center for Civic Media and Neha Narula at the Media Lab—to build a tool called Gobo that lets people aggregate and filter their feeds on their own. Why?**

EZ What this is meant to do is to say, "Look, it's really a mistake to give one or two companies control over our digital public sphere." Instead, we need competing platforms. We're trying to make the case that you want those different social networks because you want more control over your filters about what you see and what you don't see.

If we need competing platforms, we need tools that would let us use those competing platforms. Gobo is one of those tools. Gobo is an aggregator. It aggregates Twitter and the "aggregateable" parts of Facebook—the public pages.

So first we built the aggregator. And then we built the algorithms [that determine which posts you'll see]. And rather than making them a top-secret black box, we made it an open box where you can reach in and set the sliders and experiment and say, "Oh, I like how this works. Now let me change it this way and see if it works better for me."

Where we want to get in the longer term is even more of an open box; we built Gobo so that other people can write filters for it.

RM **After a lot of criticism related to the ways its news feed filters content, Facebook has started pushing posts from users' friends and family more and deemphasizing ones from brands. Do you feel this move shows Facebook is actually starting to shift its focus?**

EZ I don't believe that this is changing yet, and I won't believe it until I see a credible business model based on something other than targeted advertising.

I think that building an internet where we didn't have to pay for anything, because our attention was going to be the commodity that was traded, is one of the most destructive and shortsighted decisions that we could have made. And I do mean "we," because I was very much part of that. Until I see Facebook saying, "Look, you're going to use this as a service and you're going to pay us for the service," as opposed to "We're going to capture your attention and repackage and sell it," I won't believe it.

RM **A growing chorus of former Facebook executives and investors have been speaking out against Facebook—saying, for example, that social media is "ripping apart the fabric of how society works."**

EZ I think what's happening is that some of these people who are stepping out of the really intense "I've been in the process of building it" are starting to look at it from the outside and say, "Oh, wow, okay; now I can see the politics from the outside, and I'm not thrilled about what I've been associated with."

We need to figure out how to have those conversations a lot earlier. We should be having those conversations with people who are working at these companies and who are making these design decisions. I want to be having those conversations with my students, because my students are often going to these companies and often find themselves with the opportunity to make those design decisions.

"Because Facebook's so friggin' big, they get all sorts of advantages that make it very hard to catch up with them ... So when someone shows up as a meaningful competitor, [Facebook is] more likely to buy them and eat them up than they are to actually have to fight them in the marketplace."

RM **Why is it so hard for anyone who's not Facebook, Instagram (which is owned by Facebook), Twitter, or Snapchat to compete in this social sphere?**

EZ Network effects basically say, "I gotta be on Facebook 'cause everybody I know is on Facebook." Because Facebook's so friggin' big, they get all sorts of advantages that make it very hard to catch up with them. They get more bandwidth, they get cheaper servers.

So when someone shows up as a meaningful competitor, [Facebook is] more likely to buy them and eat them up than they are to actually have to fight them in the marketplace.

RM **You wrote a piece in the *Atlantic* that suggested a publicly supported social network as a potential solution to social media's echo-chamber effect. Could this actually happen?**

EZ I think it's wholly unrealistic in the United States. It's something that could be realistic in Europe, [where] you have a public media culture that accepts the idea that you might want to invest money in people having some basic knowledge about politics, the world, the people around them. I could imagine an innovative European public broadcaster saying, "Maybe we build a social network that's compatible with other social networks, has algorithms designed to help you tune whether you're getting news about the world, news about your community, and makes those levers visible and controllable."

TONY LUONG

Q+A

Timnit Gebru

Artificial intelligence is an increasingly seamless part of our everyday lives, present in everything from web searches to social media to home assistants like Alexa. But what do we do if this massively important technology is unintentionally, but fundamentally, biased? And what do we do if this massively important field includes almost no black researchers? Timnit Gebru is tackling these questions as part of Microsoft's Fairness, Accountability, Transparency, and Ethics in AI group, which she joined last summer. She also cofounded the Black in AI event at the Neural Information Processing Systems (NIPS) conference in 2017 and was on the steering committee for the first Fairness and Transparency conference in February. She spoke with *MIT Technology Review* about how bias gets into AI systems and how diversity can counteract it.

By Jackie Snow

JS **How does the lack of diversity distort artificial intelligence and specifically computer vision?**

TG I can talk about this for a whole year. There is a bias to what kinds of problems we think are important, what kinds of research we think are important, and where we think AI should go. If we don't have diversity in our set of researchers, we are not going to address problems that are faced by the majority of people in the world. When problems don't affect us, we don't think they're that important, and we might not even know what these problems are, because we're not interacting with the people who are experiencing them.

JS **Are there ways to counteract bias in systems?**

TG The reason diversity is really important in AI, not just in data sets but also in researchers, is that you need people who just have this social sense of how things are. We are in a diversity crisis for AI. In addition to having technical conversations, conversations about law, conversations about ethics, we need to have conversations about diversity in AI. We need all sorts of diversity in AI. And this needs to be treated as something that's extremely urgent.

From a technical standpoint, there are many different kinds of approaches. One is to diversify your data set and to have many different annotations of your data set, like race and gender and age. Once you train a model, you can test it out and see how well it does by all these different subgroups. But even after you do this, you are bound to have some sort of bias in your data set. You cannot have a data set that perfectly samples the whole world.

Something I'm really passionate about and I'm working on right now is to figure out how to encourage companies to give more information to users or even researchers. They should have recommended usage, what the pitfalls are, how biased the data set is, etc. So that when I'm a startup and I'm just taking your off-the-shelf data set or off-the-shelf model and incorporating it into whatever I'm doing, at least I have some knowledge of what kinds of pitfalls there may be. Right now we're in a place almost like the Wild West, where we don't really have many standards [about] where we put out data sets.

And then there are just some things you probably shouldn't be using machine learning for right now, and we don't have a clear guideline for what those things are. We should say that if you're going to

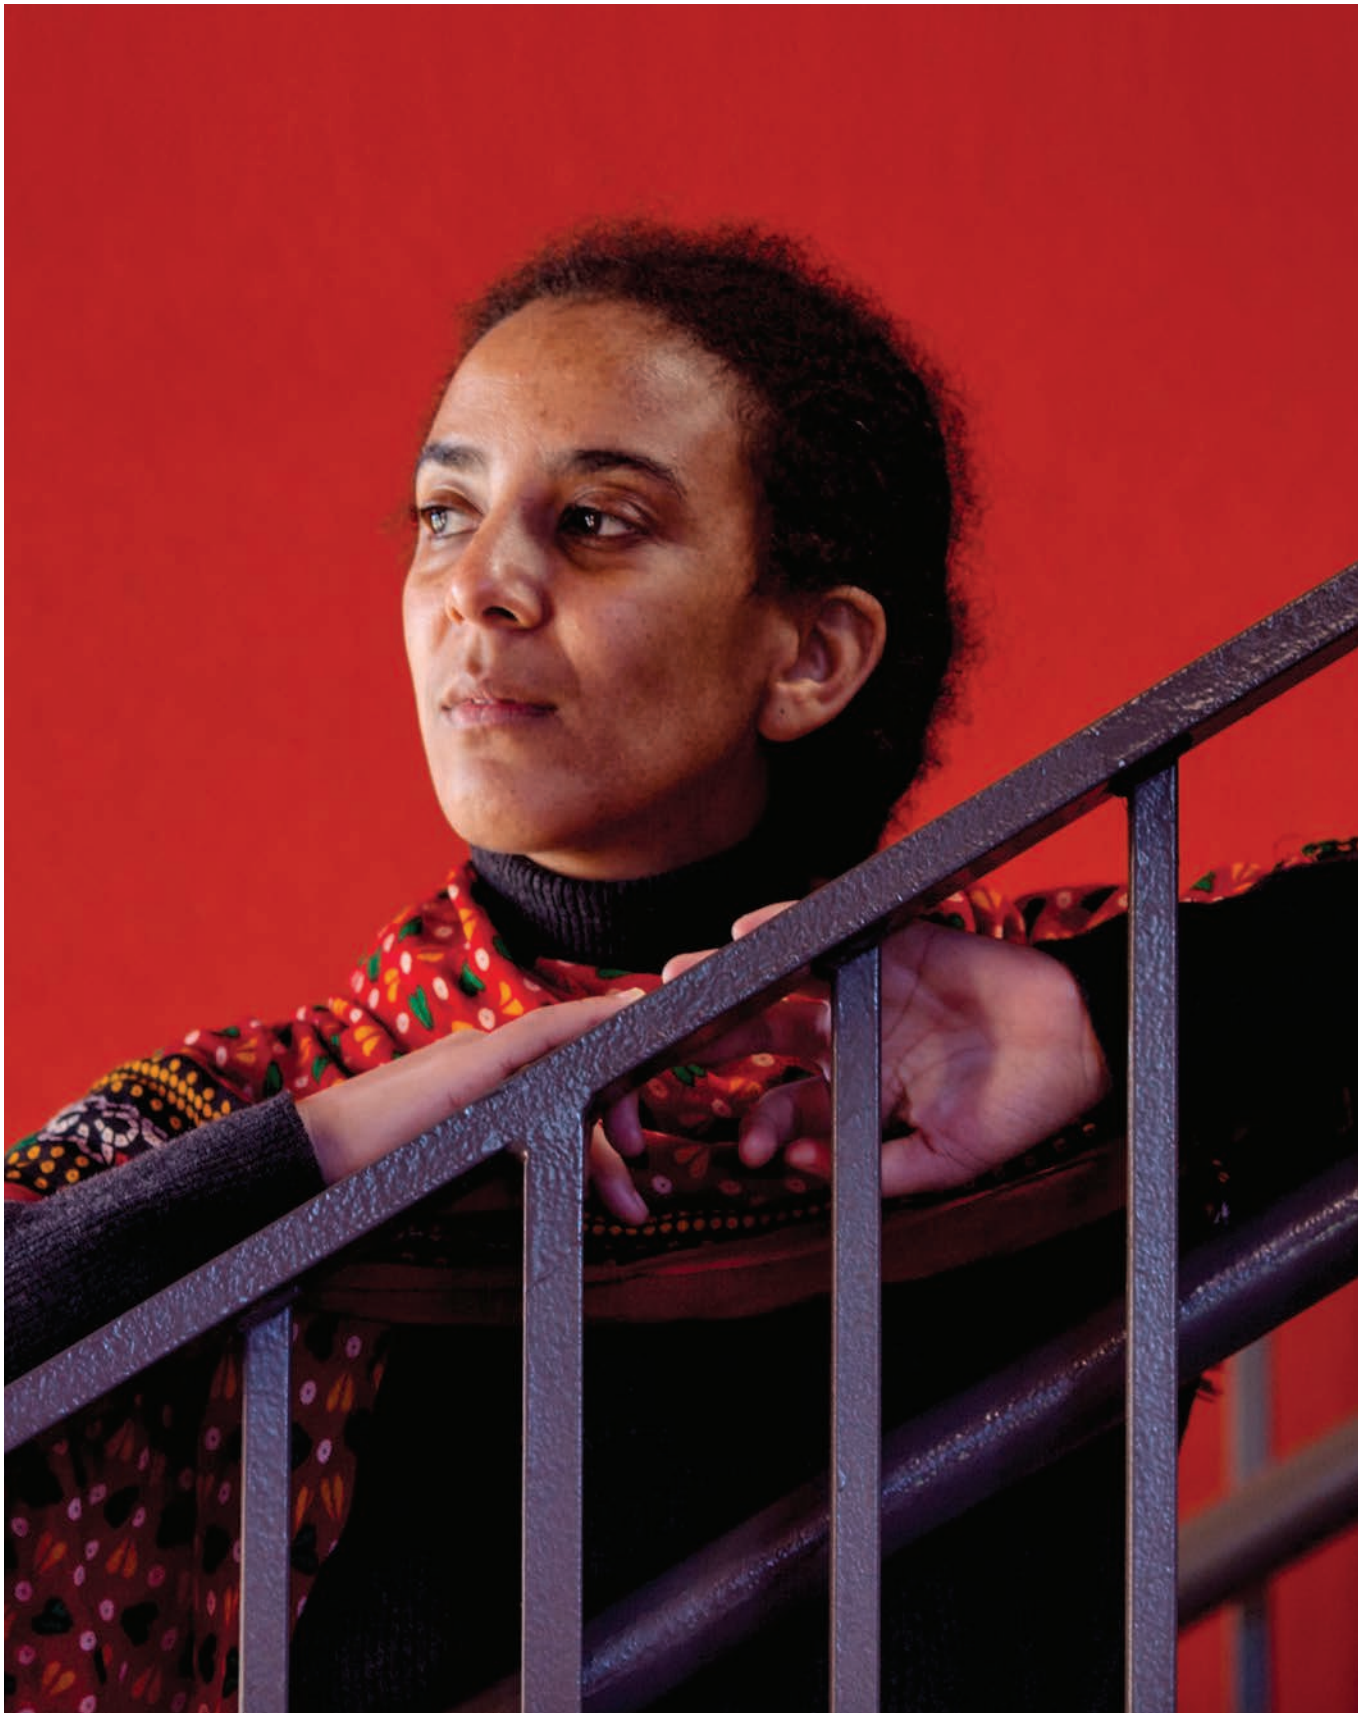

COURTESY OF TIMNIT GEBRU

COURTESY OF TIMNIT GEBRU

use machine learning for this particular task, the accuracy of your model should be at least X, and it should be fair in this particular respect. We don't have any sort of guidelines for that either. AI is just now starting to be baked into the mainstream, into a product everywhere, so we're at a precipice where we really need some sort of conversation around standardization and usage.

JS **What's been the driving motivation behind your work with Google Street View and other demographic research?**

TG At the time we started this project, there was very little work being done to try to analyze culture using images. But we know that online, most of our data is in the form of images. One of our motivations was to show that you could do social analyses using images.

This could be very useful in cases where getting survey-based data is really hard. There are places in the world where the infrastructure is not there and the resources are not there to send people door to door and gather [census] data, [but where] having an understanding of the different types of populations that live in your country would be very helpful.

But then again, this is exactly the thing that also made me want to study fairness. Because if I'm going to be continuing to do this line of work, I really need to have a better understanding of the potentially negative repercussions. What are the repercussions for surveillance? Also, what are the repercussions for a data-set bias? In any sort of data-mining project, you're going to have a bias. So my line of work there was really what led me to want to spend some time in the fairness community to understand where the pitfalls could be.

"When I started Black in AI, I started it with a couple of my friends. I had a tiny mailing list before that where I literally would add any black person I saw in this field into the mailing list and be like, 'Hi, I'm Timnit. I'm black person number two. Hi, black person number one. Let's be friends.'"

JS **What issues are you hoping to address with this first Fairness and Transparency conference?**

TG This is really the first conference that is addressing the issues of fairness, accountability, ethics, and transparency in AI. There have been workshops at other conferences, and mostly there have been workshops at either natural-language-processing-based conferences or machine-learning-based conferences. It's really important to have the stand-alone conference because it needs to be worked on by people from many disciplines who talk to each other.

Machine-learning people on their own cannot solve this problem. There are issues of transparency; there are issues of how the laws should be updated. If you're going to talk about bias in health care, you want to talk to [health-care professionals] about where the potential biases could be, and then you can think about how to have a machine-learning-based solution.

TG **What has been your experience working in AI?**

TG It's not easy. I love my job. I love the research that I work on. I love the field. I cannot imagine what else I would do in that respect. That being said, it's very difficult to be a black woman in this field. When I started Black in AI, I started it with a couple of my friends. I had a tiny mailing list before that where I literally would add any black person I saw in this field into the mailing list and be like, "Hi, I'm Timnit. I'm black person number two. Hi, black person number one. Let's be friends."

What really just made it accelerate was [in 2016] when I went to NIPS and someone was saying there were an estimated 8,500 people. I counted six black people. I was literally panicking. That's the only way I can describe how I felt. I saw that this field was growing exponentially, hitting the mainstream; it's affecting every part of society. At the same time, I also saw a lot of rhetoric about diversity and how a lot of companies think it's important.

And I saw a mismatch between the rhetoric and action. Because six black people out of 8,500—that's a ridiculous number, right? That is almost zero percent. I was like, "We have to do something now." I want to give a call to action to people who believe diversity is important. Because it is an emergency, and we have to do something about it now.

LANDON SPEERS

Q+A

Yasmin Green

Hate speech, online radicalization, bomb-making tutorials, state-sponsored fake news, web censorship—the internet is a terrible place. A portion of it is, anyway. Helping to keep that portion as small as possible is the job of Yasmin Green, director of research and development at Jigsaw, an arm of Google's parent company, Alphabet. She takes on terrorists and trolls alike with a multinational group of around 60 specialists, including software engineers, research scientists, and product managers. They can also tap into Google's huge resources. But Green and her colleagues don't spend all their time coding and building; many of them also travel to geopolitical hot spots to study radicalization up close. She spoke with *MIT Technology Review* about some of the group's methods.

By Martin Giles

MG **What drew you to the role at Jigsaw?**

YG I left Iran, where I was born, at a very young age and grew up in the UK. I remember going back at the age of 19 and feeling astounded by the level of censorship there was in the country, both offline and online. Terms like "world wide web" tend to sound ironic to people who aren't living in countries where information is free. So I was drawn to work at a company which tries to make information available to everyone.

MG **Jigsaw deals with online threats, but team members also visit conflict zones as part of their work. What's the goal of these trips?**

YG We want to make sure that we're designing technology that's based on an understanding of human experiences. A big part of our methodology is sending team members out into the field to interview people who are on the front lines of the challenges we try to tackle, whether that's repression or conflict. One of our field trips was to Iraq, where we sat face to face with people who had joined and then left the terrorist group known as ISIS. We wanted to understand the radicalization process, both the human elements but also the role that technology played. We heard about how people discovered ISIS, how people were recruited, and how technology was useful to the logistics of their travel [to join the group].

MG **What's the most important lesson we've learned about ISIS's ability to leverage the internet for recruitment?**

YG ISIS pretty much masters many media, from radio to leafleting. When it comes to the internet, they've really understood the power of microtargeting. They create content in a long list of languages, including Arabic and English, but it goes on and on, and even gets to Chinese and Hebrew. The language that really blew my mind to see a video in was sign language. So they are creating very local recruiting materials and using the algorithms that are available through social media to distribute this material and reach people in all corners of the world.

MG **Some of the content terrorist networks post online clearly needs to be taken down. But how do we deal with more subtle forms of propaganda?**

YG There are definitely categories of content that you want to make sure don't see the light of day, like beheadings and bomb-making tutorials. Then there's a whole host of other content that really isn't advocating for violence but could help advance people down the path toward it. Our research was aimed at understanding what the recruiting themes were that got people to sign up to ISIS. It turns out they weren't generally drawn to beheadings; instead, they were convinced that this group was religiously legitimate and the cause of jihad was their religious duty. ISIS was shaping the conversation, asking questions to which [it] had seemingly compelling answers.

MG **How have you tried to counter this online radicalization?**

YG One of the takeaways for us was that timing is critical. By the time potential recruits are sold on the ideology, it's too late to influence them; you have to get to them when they are sympathetic but not yet sold. So we turned to targeted online advertising. We've designed something called the "redirect method," which uses targeted ads to reach people who are sympathetic to ISIS but not yet committed to it, and redirects them to online videos from moderate clerics, defectors, and citizen journalists that could avert their radicalization.

MG **What results has this method delivered?**

YG The pilot, which was eight weeks long and ran in Arabic and English, reached 320,000 people. It had an exceptionally high click-through rate on the ads and drove half a million minutes of video watch time. Given people don't spend more than a few seconds on a video they're not interested in, that's encouraging. After our pilot, YouTube integrated the method into its search results. The open-source methodology has also been replicated by others like the Gen Next Foundation, and we continue to support some new deployments.

MG **Jigsaw also tries to tackle online censorship. How bad is this problem?**

YG If you look at Freedom House's index on this, they say that every year the situation is getting worse. That's really discouraging, because it's antithetical to what people who are developing the internet want for it. The situation is so volatile. When you have civil unrest, as there was recently in Iran, you see the dilemma of repressive governments around whether or not to shut down the internet because [censoring] it inflames the population and draws public attention and outrage. But if those who are in power feel threatened, they will censor.

MG **What can companies like Alphabet do to counter this?**

YG We feel a really big responsibility to help people get access to information, especially when there's conflict and repression. One of our products focuses on protecting independent media around the world from a type of censorship attack called distributed denial of service, or DDoS, which knocks websites offline by flooding them with traffic. It's called Project Shield. The idea for this came from our fieldwork. Our team spoke to the Kenyan election monitoring group, and their site had gone down on the day of a key election. Google has an enormous infrastructure and world-class DDoS mitigation capabilities, which Project Shield takes advantage of. Without websites having to host with Google or Jigsaw, it offers to vet traffic before it arrives at a server so we can spot the traffic that's malicious and filter it out.

MG **Let's switch to the problem of fake news. How can we better identify state-sponsored disinformation efforts?**

YG The goal of these campaigns is to plant ideas and narratives via fake personas and news sites, and then have these seeds fertilized by the masses so that the conversations look like they're organic. The research question we've been asking ourselves is whether there are technical markers of coordinated, covert activity versus the more organic activity we expect to see online. If there are, we could use these in automated detection.

We're looking at several dimensions here. Coordinated campaigns tend to outlast organic ones, and the people involved tend to wait to get instructions from their masters, so there's sometimes a slight delay before they suddenly act together. Those actors also tend to be linked in extremely tight networks or clusters that look anomalous. There's also a semantic dimension. These campaigns often use similar words and phrases that can be a signal of centralized control.

MG **Web companies have been criticized for not doing enough to address online hate speech and harassment. Is that criticism fair?**

YG I don't think anyone imagined that we'd have the level of intimidation and hate speech online that we currently do, and we have to try really hard to make sure that we're getting ahead of it.

MG **What is Jigsaw doing to frustrate online trolls?**

YG We have a team dedicated to seeing how natural-language processing and machine learning can be used to identify online toxicity and help moderators and communities in tackling it. There's a model we've developed that's publicly available, called Perspective, and you can find it at www.perspectiveapi.com. This helps you score comments for their level of toxicity. The research team is looking at ways to get to another level of granularity to help us better identify what's happening and how moderators can control it.

MG **There's been a lot of concern that bias in algorithms could harm certain groups in society. Is this something you're also focused on?**

YG Yes. We're spending a lot of time thinking about how we can ensure that we don't have biases in our AI in Perspective that could be harmful to the goals of that project, which is to create inclusive and empathetic conversations.

MG **You spend a lot of your life focused on the dark side of the internet. Are you still optimistic about its potential?**

YG I am, but we need to keep developing innovative technologies that can help address these really hard challenges.

“I don’t think anyone imagined that we’d have the level of intimidation and hate speech online that we currently do, and we have to try really hard to make sure that we’re getting ahead of it.”

LANDON SPEERS

10 BREAK TECHNOL

Every year since 2001 we've picked what we call the 10 Breakthrough Technologies. People often ask, what exactly do you mean by "breakthrough"? It's a reasonable question—some of our picks haven't yet reached widespread use, while others may be on the cusp of becoming commercially available. What we're really looking for is a technology, or perhaps even a collection of technologies, that will have a profound effect on our lives. For this year, a new technique in artificial intelligence called GANs is giving machines imagination; artificial embryos, despite some thorny ethical constraints, are redefining how life can be created and are opening a research window into the early moments of a human life; and a pilot plant in the heart of Texas's petrochemical industry is attempting to create completely clean power from natural gas—probably a major energy source for the foreseeable future. These and the rest of our list will be worth keeping an eye on.

—The Editors

KTHROUGH
_____OGIES

The 2018 list p38

The GANfather: the man who's given machines the gift of imagination p48

Forecasts of genomic fate just got a lot more accurate p54

This is how a Google sister company envisions the future of cities p60

Quantum computers are finally here. What are we going to do with them? p66

3-D METAL PRINTING

New machines are making 3-D printing of metal parts practical for the first time.

BREAKTHROUGH
Now printers can make metal objects quickly and cheaply.

WHY IT MATTERS
The ability to make large and complex metal objects on demand could transform manufacturing.

KEY PLAYERS
Markforged
Desktop Metal
GE

AVAILABILITY
Now

While 3-D printing has been around for decades, it has remained largely in the domain of hobbyists and designers producing one-off prototypes. And printing objects with anything other than plastics—in particular, metal—has been expensive and painfully slow.

Now, however, it's becoming cheap and easy enough to be a potentially practical way of manufacturing parts. If widely adopted, it could change the way we mass-produce many products.

In the short term, manufacturers wouldn't need to maintain large inventories—they could simply print an object, such as a replacement part for an aging car, whenever someone needs it.

In the longer term, large factories that mass-produce a limited range of parts might be replaced

DEREK BRAHNEY

by smaller ones that make a wider variety, adapting to customers' changing needs.

The technology can create lighter, stronger parts, and complex shapes that aren't possible with conventional metal fabrication methods. It can also provide more precise control of the microstructure of metals. In 2017, researchers from the Lawrence Livermore National Laboratory announced they had developed a 3-D-printing method for creating stainless-steel parts twice as strong as traditionally made ones.

Also in 2017, 3-D-printing company Markforged, a small startup based outside Boston, released the first 3-D metal printer for under $100,000.

Another Boston-area startup, Desktop Metal, began to ship its first metal prototyping machines in December 2017. It plans to begin selling larger machines, designed for manufacturing, that are 100 times faster than older metal printing methods.

The printing of metal parts is also getting easier. Desktop Metal now offers software that generates designs ready for 3-D printing. Users tell the program the specs of the object they want to print, and the software produces a computer model suitable for printing.

GE, which has long been a proponent of using 3-D printing in its aviation products (see "10 Breakthrough Technologies of 2013: Additive Manufacturing"), has a test version of its new metal printer that is fast enough to make large parts. The company plans to begin selling the printer in 2018.

—Erin Winick

ARTIFICIAL EMBRYOS

Scientists have begun to forge embryos out of stem cells.

BREAKTHROUGH
Without using eggs or sperm cells, researchers have made embryo-like structures from stem cells alone, providing a whole new route to creating life.

WHY IT MATTERS
Artificial embryos will make it easier for researchers to study the mysterious beginnings of a human life, but they're stoking new bioethical debates.

KEY PLAYERS
University of Cambridge
University of Michigan
Rockefeller University

AVAILABILITY
Now

In a breakthrough that redefines how life can be created, embryologists working at the University of Cambridge in the UK have grown realistic-looking mouse embryos using only stem cells. No egg. No sperm. Just cells plucked from another embryo.

The researchers placed the cells carefully in a three-dimensional scaffold and watched, fascinated, as they started communicating and lining up into the distinctive bullet shape of a mouse embryo several days old.

"We know that stem cells are magical in their powerful potential of what they can do. We did not realize they could self-organize so beautifully or perfectly," Magdelena Zernicka-Goetz, who headed the team, told an interviewer at the time.

Zernicka-Goetz says her "synthetic" embryos probably couldn't have grown into mice. Nonetheless, they're a hint that soon we could have mammals born without an egg at all.

That isn't Zernicka-Goetz's goal. She wants to study how the cells of an early embryo begin taking on their specialized roles. The next step, she says, is to make an artificial embryo out of human stem cells, work that's being pursued at the University of Michigan and Rockefeller University.

Synthetic human embryos would be a boon to scientists, letting them tease apart events early in development. And since such embryos start with easily manipulated stem cells, labs will be able to employ a full range of tools, such as gene editing, to investigate them as they grow.

Artificial embryos, however, pose ethical questions. What if they turn out to be indistinguishable from real embryos? How long can they be grown in the lab before they feel pain? We need to address those questions before the science races ahead much further, bioethicists say.

—Antonio Regalado

SENSING CITY

Alphabet's Sidewalk Labs plans to create a high-tech district to rethink how we build and run cities.

BREAKTHROUGH
A Toronto neighborhood aims to be the first place to successfully integrate cutting-edge urban design with state-of-the-art digital technology.

WHY IT MATTERS
Smart cities could make urban areas more affordable, livable, and environmentally friendly.

AVAILABILITY
Project announced in October 2017; construction could begin in 2019

KEY PLAYERS
Sidewalk Labs
Waterfront Toronto

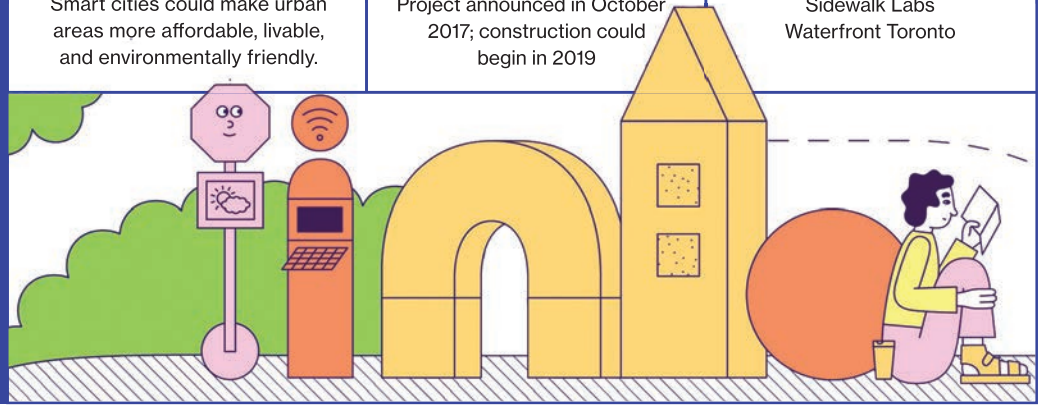

AI FOR EVERYBODY

Making machine-learning tools available through cloud services could spread artificial intelligence far and wide.

BREAKTHROUGH
Cloud-based AI is making the technology cheaper and easier to use.

KEY PLAYERS
Amazon
Google
Microsoft

WHY IT MATTERS
Right now the use of AI is dominated by a relatively few companies, but as a cloud-based service, it could be widely available to many more, giving the economy a boost.

AVAILABILITY
Now

Artificial intelligence has so far been mainly the plaything of big tech companies like Amazon, Baidu, Google, and Microsoft, as well as some startups. For many other companies and parts of the economy, AI systems are too expensive and too difficult to implement fully.

What's the solution? Machine-learning tools based in the cloud are bringing AI to a far broader audience. So far, Amazon dominates cloud AI with its AWS subsidiary. Google is challenging that with TensorFlow, an open-source AI library that can be used to build other machine-learning software. Recently Google announced Cloud AutoML, a suite of pre-trained systems that could make AI simpler to use.

Microsoft, which has its own AI-powered cloud platform, Azure, is teaming up with Amazon to offer Gluon, an open-source deep-learning library. Gluon is supposed to make building neural nets—a key technology in AI that crudely mimics how the human brain learns—as easy as building a smartphone app.

It is uncertain which of these companies will become the leader in offering AI cloud services. But it is a huge business opportunity for the winners.

These products will be essential if the AI revolution is going to

MARTINA PAUKOVA

Numerous smart-city schemes have run into delays, dialed down their ambitious goals, or priced out everyone except the super-wealthy. A new project in Toronto, called Quayside, is hoping to change that pattern of failures by rethinking an urban neighborhood from the ground up and rebuilding it around the latest digital technologies.

Alphabet's Sidewalk Labs, based in New York City, is collaborating with the Canadian government on the high-tech project, slated for Toronto's industrial waterfront.

One of the project's goals is to base decisions about design, policy, and technology on information from an extensive network of sensors that gather data on everything from air quality to noise levels to people's activities.

The plan calls for all vehicles to be autonomous and shared. Robots will roam underground doing menial chores like delivering the mail. Sidewalk Labs says it will open access to the software and systems it's creating so other companies can build services on top of them, much as people build apps for mobile phones.

The company intends to closely monitor public infrastructure, and this has raised concerns about data governance and privacy. But Sidewalk Labs believes it can work with the community and the local government to alleviate those worries.

"What's distinctive about what we're trying to do in Quayside is that the project is not only extraordinarily ambitious but also has a certain amount of humility," says Rit Aggarwala, the executive in charge of Sidewalk Labs' urban-systems planning. That humility may help Quayside avoid the pitfalls that have plagued previous smart-city initiatives.

Other North American cities are already clamoring to be next on Sidewalk Labs' list, according to Waterfront Toronto, the public agency overseeing Quayside's development. "San Francisco, Denver, Los Angeles, and Boston have all called asking for introductions," says the agency's CEO, Will Fleissig.

—Elizabeth Woyke

FEATURE ON PAGE 60

spread more broadly through different parts of the economy.

Currently AI is used mostly in the tech industry, where it has created efficiencies and produced new products and services. But many other businesses and industries have struggled to take advantage of the advances in artificial intelligence. Sectors such as medicine, manufacturing, and energy could also be transformed if they were able to implement the technology more fully, with a huge boost to economic productivity.

Most companies, though, still don't have enough people who know how to use cloud AI. So Amazon and Google are also setting up consultancy services. Once the cloud puts the technology within the reach of almost everyone, the real AI revolution can begin.

—Jackie Snow

MIGUEL PORLAN

DUELING NEURAL NETWORKS

By playing cat-and-mouse games with data, a pair of AI systems can acquire an imagination.

BREAKTHROUGH
Two AI systems can spar with each other to create ultra-realistic original images or sounds, something machines have never been able to do before.

WHY IT MATTERS
This gives machines something akin to a sense of imagination, which may help them become less reliant on humans—but also turns them into alarmingly powerful tools for digital fakery.

KEY PLAYERS
Google Brain, DeepMind, Nvidia

AVAILABILITY
Now

Artificial intelligence is getting very good at identifying things: show it a million pictures, and it can tell you with uncanny accuracy which ones depict a pedestrian crossing a street. But AI is hopeless at generating images of pedestrians by itself. If it could do that, it would be able to create gobs of realistic but synthetic pictures depicting pedestrians in various settings, which a self-driving car could use to train itself without ever going out on the road.

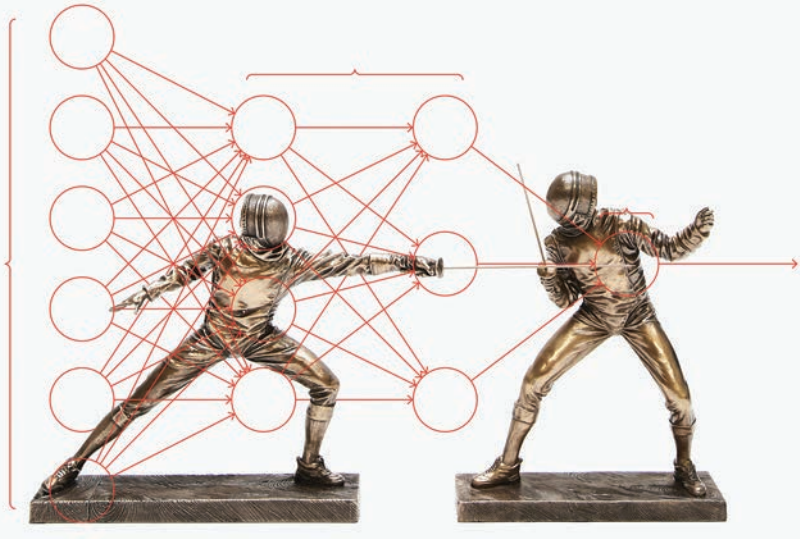

DEREK BRAHNEY

The problem is, creating something entirely new requires imagination—and until now that has perplexed AIs.

The solution first occurred to Ian Goodfellow, then a PhD student at the University of Montreal, during an academic argument in a bar in 2014. The approach, known as a generative adversarial network, or GAN, takes two neural networks—the simplified mathematical models of the human brain that underpin most modern machine learning—and pits them against each other in a digital cat-and-mouse game.

Both networks are trained on the same data set. One, known as the generator, is tasked with creating variations on images it's already seen—perhaps a picture of a pedestrian with an extra arm. The second, known as the discriminator, is asked to identify whether the example it

sees is like the images it has been trained on or a fake produced by the generator—basically, is that three-armed person likely to be real?

Over time, the generator can become so good at producing images that the discriminator can't spot fakes. Essentially, the generator has been taught to recognize, and then create, realistic-looking images of pedestrians.

The technology has become one of the most promising advances in AI in the past decade, able to help machines produce results that fool even humans.

GANs have been put to use creating realistic-sounding speech and photorealistic fake imagery. In one compelling example, researchers from chipmaker Nvidia primed a GAN with celebrity photographs to create hundreds of credible faces of people who don't exist. Another research group made not-unconvincing fake paintings that look like the works of van Gogh. Pushed further, GANs can reimagine images in different ways—making a sunny road appear snowy, or turning horses into zebras.

The results aren't always perfect: GANs can conjure up bicycles with two sets of handlebars, say, or faces with eyebrows in the wrong place. But because the images and sounds are often startlingly realistic, some experts believe there's a sense in which GANs are beginning to understand the underlying structure of the world they see and hear. And that means AI may gain, along with a sense of imagination, a more independent ability to make sense of what it sees in the world.

—Jamie Condliffe

FEATURE ON PAGE 48

BABEL-FISH EARBUDS

Google's Pixel Buds show the promise of real-time translation, though the current hardware is clunky.

BREAKTHROUGH
Near-real-time translation now works for a large number of languages and is easy to use.

WHY IT MATTERS
In an increasingly global world, language is still a barrier to communication.

KEY PLAYERS
Google
Baidu

AVAILABILITY
Now

In the cult sci-fi classic *The Hitchhiker's Guide to the Galaxy,* you slide a yellow Babel fish into your ear to get translations in an instant. In the real world, Google has come up with an interim solution: a $159 pair of earbuds, called Pixel Buds. These work with its Pixel smartphones and Google Translate app to produce practically real-time translation.

One person wears the earbuds, while the other holds a phone. The earbud wearer speaks in his or her language—English is the default—and the app translates the talking and plays it aloud on the phone. The person holding the phone responds; this response is translated and played through the earbuds.

Google Translate already has a conversation feature, and its iOS and Android apps let two users speak as it automatically figures out what languages they're using and then translates them. But background noise can make it hard for the app to understand what people are saying, and also to figure out when one person has stopped speaking and it's time to start translating.

Pixel Buds get around these problems because the wearer taps and holds a finger on the right earbud while talking. Splitting the interaction between the phone and the earbuds gives each person control of a microphone and helps the speakers maintain eye contact, since they're not trying to pass a phone back and forth.

The Pixel Buds were widely panned for subpar design. They do look silly, and they may not fit well in your ears. They can also be hard to set up with a phone.

Clunky hardware can be fixed, though. Pixel Buds show the promise of mutually intelligible communication between languages in close to real time. And no fish required.

—Rachel Metz

ZERO CARBON NATURAL GAS

A new engineering approach to natural-gas plants puts carbon dioxide to work.

BREAKTHROUGH
A power plant efficiently and cheaply captures carbon released by burning natural gas, avoiding greenhouse-gas emissions.

WHY IT MATTERS
Around 32 percent of US electricity is produced with natural gas, accounting for around 30 percent of the power sector's carbon emissions.

KEY PLAYERS
8 Rivers Capital
Exelon Generation
CB&I

AVAILABILITY
3 to 5 years

The world is probably stuck with natural gas as one of our primary sources of electricity for the foreseeable future. Cheap and readily available, it now accounts for more than 30 percent of US electricity and 22 percent of world electricity. And although it's cleaner than coal, it's still a massive source of carbon emissions.

A pilot power plant just outside Houston, in the heart of the US petroleum and refining industry, is testing a technology that could make clean energy from natural gas a reality. The company behind the 50-megawatt project, Net Power, believes it can generate power at least as cheaply as standard natural-gas plants and capture essentially all the carbon dioxide released in the process.

If so, it would mean the world has a way to produce carbon-free energy from a fossil fuel at a reasonable cost. Such natural-gas plants could be cranked up and down on demand, avoiding the high capital costs of nuclear power and sidestepping the unsteady supply that renewables generally provide.

Net Power is a collaboration between technology development firm 8 Rivers Capital, Exelon Generation, and energy construction firm CB&I. The company is in the process of commissioning the plant and has begun initial testing. It intends to release results from early evaluations in the months ahead.

The plant puts the carbon dioxide released from burning natural gas under high pressure and heat, using the resulting supercritical CO_2 as the "working fluid" that drives a specially built turbine. Much of the carbon dioxide can be continuously recycled; the rest can be captured cheaply.

A key part of pushing down the costs depends on selling that carbon dioxide. Today the main use is in helping to extract oil from petroleum wells. That's a limited market, and not a particularly green one. Eventually, however, Net Power hopes to see growing demand for carbon dioxide in cement manufacturing and in making plastics and other carbon-based materials.

Net Power's technology won't solve all the problems with natural gas, particularly on the extraction side. But as long as we're using natural gas, we might as well use it as cleanly as possible. Of all the clean-energy technologies in development, Net Power's is one of the furthest along to promise more than a marginal advance in cutting carbon emissions.

—James Temple

MIGUEL PORLAN

PERFECT ONLINE PRIVACY

A tool developed for blockchains makes it possible to carry out a digital transaction without revealing any more information than absolutely necessary.

BREAKTHROUGH
Computer scientists are perfecting a cryptographic tool for proving something without revealing the information underlying the proof.

WHY IT MATTERS
If you need to disclose personal information to get something done online, it will be easier to do so without risking your privacy or exposing yourself to identity theft.

KEY PLAYERS
Zcash
JPMorgan Chase
ING

AVAILABILITY
Now

True internet privacy could finally become possible thanks to a new tool that can—for instance—let you prove you're over 18 without revealing your date of birth, or prove you have enough money in the bank for a financial transaction without revealing your balance or other details. That limits the risk of a privacy breach or identity theft.

The tool is an emerging cryptographic protocol called a zero-knowledge proof. Though researchers have worked on it for decades, interest has exploded in the past year, thanks in part to the growing obsession with cryptocurrencies, most of which aren't private.

Much of the credit for a practical zero-knowledge proof goes to Zcash, a digital currency that launched in late 2016. Zcash's developers used a method called a zk-SNARK (for "zero-knowledge succinct non-interactive argument of knowledge") to give users the power to transact anonymously.

That's not normally possible in Bitcoin and most other public blockchain systems, in which transactions are visible to everyone. Though these transactions are theoretically anonymous, they can be combined with other data to track and even identify users. Vitalik Buterin, creator of Ethereum, the world's second-most-popular blockchain network, has described zk-SNARKs as an "absolutely game-changing technology."

For banks, this could be a way to use blockchains in payment systems without sacrificing their clients' privacy. Last year, JPMorgan Chase added zk-SNARKs to its own blockchain-based payment system.

For all their promise, though, zk-SNARKs are computation-heavy and slow. They also require a so-called "trusted setup," creating a cryptographic key that could compromise the whole system if it fell into the wrong hands. But researchers are looking at alternatives that deploy zero-knowledge proofs more efficiently and don't require such a key.

—Mike Orcutt

GENETIC FORTUNE-TELLING

Large genetic studies are allowing scientists to predict common diseases and human traits.

BREAKTHROUGH
Scientists can now use your genome to predict your chances of getting heart disease or breast cancer, and even your IQ.

WHY IT MATTERS
DNA-based predictions could be the next great public health advance, but they will increase the risks of genetic discrimination.

KEY PLAYERS
Helix
23andMe
Myriad Genetics
UK Biobank
Broad Institute

AVAILABILITY
Now

One day, babies will get DNA report cards at birth. These reports will offer predictions about their chances of suffering a heart attack or cancer, of getting hooked on tobacco, and of being smarter than average.

The science making these report cards possible has suddenly arrived, thanks to huge genetic studies—some involving more than a million people.

It turns out that most common diseases and many behaviors and traits, including intelligence, are a result of not one or a few genes but many acting in concert. Using the

DEREK BRAHNEY

data from large ongoing genetic studies, scientists are creating what they call "polygenic risk scores."

Though the new DNA tests offer probabilities, not diagnoses, they could greatly benefit medicine. For example, if women at high risk for breast cancer got more mammograms and those at low risk got fewer, those exams might catch more real cancers and set off fewer false alarms.

Pharmaceutical companies can also use the scores in clinical trials of preventive drugs for such illnesses as Alzheimer's or heart disease. By picking volunteers who are more likely to get sick, they can more accurately test how well the drugs work.

The trouble is, the predictions are far from perfect. Who wants to know they *might* develop Alzheimer's? What if someone with a low risk score for cancer puts off being screened, and then develops cancer anyway?

Polygenic scores are also controversial because they can predict any trait, not only diseases. For instance, they can now forecast about 10 percent of a person's performance on IQ tests. As the scores improve, it's likely that DNA IQ predictions will become routinely available. But how will parents and educators use that information?

To behavioral geneticist Eric Turkheimer, the chance that genetic data will be used for both good and bad is what makes the new technology "simultaneously exciting and alarming."

—Antonio Regalado

FEATURE ON PAGE 54

MATERIALS' QUANTUM LEAP

Researchers recently used a quantum computer to model a simple molecule. That's just the start.

BREAKTHROUGH
IBM has simulated the electronic structure of a small molecule, using a seven-qubit quantum computer.

WHY IT MATTERS
Understanding molecules in exact detail will allow chemists to design more effective drugs and better materials for generating and distributing energy.

KEY PLAYERS
IBM
Google
Harvard's Alán Aspuru-Guzik

AVAILABILITY
5 to 10 years

The prospect of powerful new quantum computers comes with a puzzle. They'll be capable of feats of computation inconceivable with today's machines, but we haven't yet figured out what we might do with those powers.

One likely and enticing possibility: precisely designing molecules.

Chemists are already dreaming of new proteins for far more effective drugs, novel electrolytes for better batteries, compounds that could turn sunlight directly into a liquid fuel, and much more efficient solar cells.

We don't have these things because molecules are ridiculously hard to model on a classical computer. Try simulating the behavior of the electrons in even a relatively simple molecule and you run into complexities far beyond the capabilities of today's computers.

But it's a natural problem for quantum computers, which instead of digital bits representing *1*s and *0*s use "qubits" that are themselves quantum systems. Recently, IBM researchers used a quantum computer with seven qubits to model a small molecule made of three atoms.

It should become possible to accurately simulate far larger and more interesting molecules as scientists build machines with more qubits and, just as important, better quantum algorithms. *—David Rotman*

FEATURE ON PAGE 66

The **GAN**father:

The man who's given machines the gift of imagination

By pitting neural networks against one another, Ian Goodfellow has created a powerful AI tool. Now he, and the rest of us, must face the consequences.

by MARTIN GILES Photography by CHRISTIE HEMM KLOK

One night in 2014, Ian Goodfellow went drinking to celebrate with a fellow doctoral student who had just graduated. At Les 3 Brasseurs (The Three Brewers), a favorite Montreal watering hole, some friends asked for his help with a thorny project they were working on: a computer that could create photos by itself.

Researchers were already using neural networks, algorithms loosely modeled on the web of neurons in the human brain, as "generative" models to create plausible new data of their own. But the results were often not very good: images of a computer-generated face tended to be blurry or have errors like missing ears. The plan Goodfellow's friends were proposing was to use a complex statistical analysis of the elements that make up a photograph to help machines come up with images by themselves. This would have required a massive amount of number-crunching, and Goodfellow told them it simply wasn't going to work.

But as he pondered the problem over his beer, he hit on an idea. What if you pitted two neural networks against each other? His friends were skeptical, so once he got home, where his girlfriend was already fast asleep, he decided to give it a try. Goodfellow coded into the early hours and then tested his software. It worked the first time.

What he invented that night is now called a GAN, or "generative adversarial network." The technique has sparked huge excitement in the field of machine learning and turned its creator into an AI celebrity.

In the last few years, AI researchers have made impressive progress using a technique called deep learning. Supply a deep-learning system with enough images and it learns to, say, recognize a pedestrian who's about to cross a road. This approach has made possible things like self-driving cars and the conversational technology that powers Alexa, Siri, and other virtual assistants.

In the future, computers will get much better at feasting on raw data and working out what they need to learn from it.

But while deep-learning AIs can learn to recognize things, they have not been good at creating them. The goal of GANs is to give machines something akin to an imagination.

Doing so wouldn't merely enable them to draw pretty pictures or compose music; it would make them less reliant on humans to instruct them about the world and the way it works. Today, AI programmers often need to tell a machine exactly what's in the training data it's being fed—which of a million pictures contain a pedestrian crossing a road, and which don't. This is not only costly and labor-intensive; it limits how well the system deals with even slight departures from what it was trained on. In the future, computers will get much better at feasting on raw data and working out what they need to learn from it without being told.

That will mark a big leap forward in what's known in AI as "unsupervised learning." A self-driving car could teach itself about many different road conditions without leaving the garage. A robot could anticipate the obstacles it might encounter in a busy warehouse without needing to be taken around it.

That will mark a big leap forward in what is known in AI as "unsupervised learning."

Our ability to imagine and reflect on many different scenarios is part of what makes us human. And when future historians of technology look back, they're likely to see GANs as a big step toward creating machines with a human-like consciousness. Yann LeCun, Facebook's chief AI scientist, has called GANs "the coolest idea in deep learning in the last 20 years." Another AI luminary, Andrew Ng, the former chief scientist of China's Baidu, says GANs represent "a significant and fundamental advance" that's inspired a growing global community of researchers.

THE GANFATHER, PART II: AI FIGHT CLUB

Goodfellow is now a research scientist on the Google Brain team, at the company's headquarters in Mountain View, California. When I met him there recently, he still seemed surprised by his superstar status, calling it "a little surreal." Perhaps no less surprising is that, having made his discovery, he now spends much of his time working against those who wish to use it for evil ends.

The magic of GANs lies in the rivalry between the two neural nets. It mimics the back-

and-forth between a picture forger and an art detective who repeatedly try to outwit one another. Both networks are trained on the same data set. The first one, known as the generator, is charged with producing artificial outputs, such as photos or handwriting, that are as realistic as possible. The second, known as the discriminator, compares these with genuine images from the original data set and tries to determine which are real and which are fake. On the basis of those results, the generator adjusts its parameters for creating new images. And so it goes, until the discriminator can no longer tell what's genuine and what's bogus.

In one widely publicized example last year, researchers at Nvidia, a chip company heavily invested in AI, trained a GAN to generate pictures of imaginary celebrities by studying real ones. Not all the fake stars it produced were perfect, but some were impressively realistic. Unlike other machine-learning approaches that require tens of thousands of training images, GANs can become proficient with a few hundred.

"There are a lot of areas of science and engineering where we need to optimize something.

This power of imagination is still limited. Once it's been trained on a lot of dog photos, a GAN can generate a convincing fake image of a dog that has, say, a different pattern of spots; but it can't conceive of an entirely new animal. The quality of the original training data also has a big influence on the results. In one telling example, a GAN began producing pictures of cats with random letters integrated into the images. Because the training data contained cat memes from the internet, the machine had taught itself that words were part of what it meant to be a cat.

GANs are also temperamental, says Pedro Domingos, a machine-learning researcher at the University of Washington. If the discriminator is too easy to fool, the generator's output won't look realistic. And calibrating the two dueling neural nets can be difficult, which explains why GANs sometimes spit out bizarre stuff such as animals with two heads.

Still, the challenges haven't deterred researchers. Since Goodfellow and a few others published the first study on his discovery, in 2014, hundreds of GAN-related papers have been written. One fan of the technology has even created a web page called the "GAN zoo," dedicated to keeping track of the various versions of the technique that have been developed.

The most obvious immediate applications are in areas that involve a lot of imagery, such as video games and fashion: what, for instance, might a game character look like running through the rain? But looking ahead, Goodfellow thinks GANs will drive more significant advances. "There are a lot of areas of science and engineering where we need to optimize something," he says, citing examples such as medicines that need to be more effective or batteries that must get more efficient. "That's going to be the next big wave."

In high-energy physics, scientists use powerful computers to simulate the likely interactions of hundreds of subatomic particles in machines like the Large Hadron Collider at CERN in Switzerland. These simulations are slow and require massive computing power. Researchers at Yale University and Lawrence Berkeley National Laboratory have developed a GAN that, after training on existing simulation data, learns to generate pretty accurate predictions of how a particular particle will behave, and does it much faster.

Medical research is another promising field. Privacy concerns mean researchers sometimes can't get enough real patient data to, say, analyze why a drug didn't work. GANs can help solve this problem by generating fake records that are almost as good as the real thing, says Casey

That's going to be the next big wave."

A GAN trained on photos of real celebrities came up with its own set of imaginary stars. In most cases, the fakes looked pretty realistic.

NVIDIA

Getting GANS to work well can be tricky. If there are glitches, the results can be bizarre.

Greene of the University of Pennsylvania. This data could be shared more widely, helping to advance research, while the real records are tightly protected.

THE GANFATHER, PART III: THE BAD FELLOWS

There is a darker side, however. A machine designed to create realistic fakes is a perfect weapon for purveyors of fake news who want to influence everything from stock prices to elections. AI tools are already being used to put pictures of other people's faces on the bodies of porn stars and put words in the mouths of politicians. GANs didn't create this problem, but they'll make it worse.

Hany Farid, who studies digital forensics at Dartmouth College, is working on better ways to spot fake videos, such as detecting slight changes in the color of faces caused by inhaling and exhaling that GANs find hard to mimic precisely. But he warns that GANs will adapt in turn. "We're fundamentally in a weak position," says Farid.

This cat-and-mouse game will play out in cybersecurity, too. Researchers are already highlighting the risk of "black box" attacks, in which GANs are used to figure out the machine-learning models with which plenty of security programs spot malware. Having divined how a defender's algorithm works, an attacker can evade it and insert rogue code. The same approach could also be used to dodge spam filters and other defenses.

Goodfellow is well aware of the dangers. Now heading a team at Google that's focused on making machine learning more secure, he warns that the AI community must learn the lesson of previous waves of innovation, in which technologists treated security and privacy as an afterthought. By the time they woke up to the risks, the bad guys had a significant lead. "Clearly, we're already beyond the start," he says, "but hopefully we can make significant advances in security before we're too far in."

Nonetheless, he doesn't think there will be a purely technological solution to fakery. Instead, he believes, we'll have to rely on societal ones, such as teaching kids critical thinking by getting them to take things like speech and debating classes. "In speech and debate you're competing against another student," he says, "and you're thinking about how to craft misleading claims, or how to craft correct claims that are very persuasive." He may well be right, but his conclusion that technology can't cure the fake-news problem is not one many will want to hear.

Goodfellow's creation can be used to imagine all sorts of things, including new interior designs.

Martin Giles is MIT Technology Review*'s San Francisco bureau chief.*

ALEC RADFORD / IAN GOODFELLOW

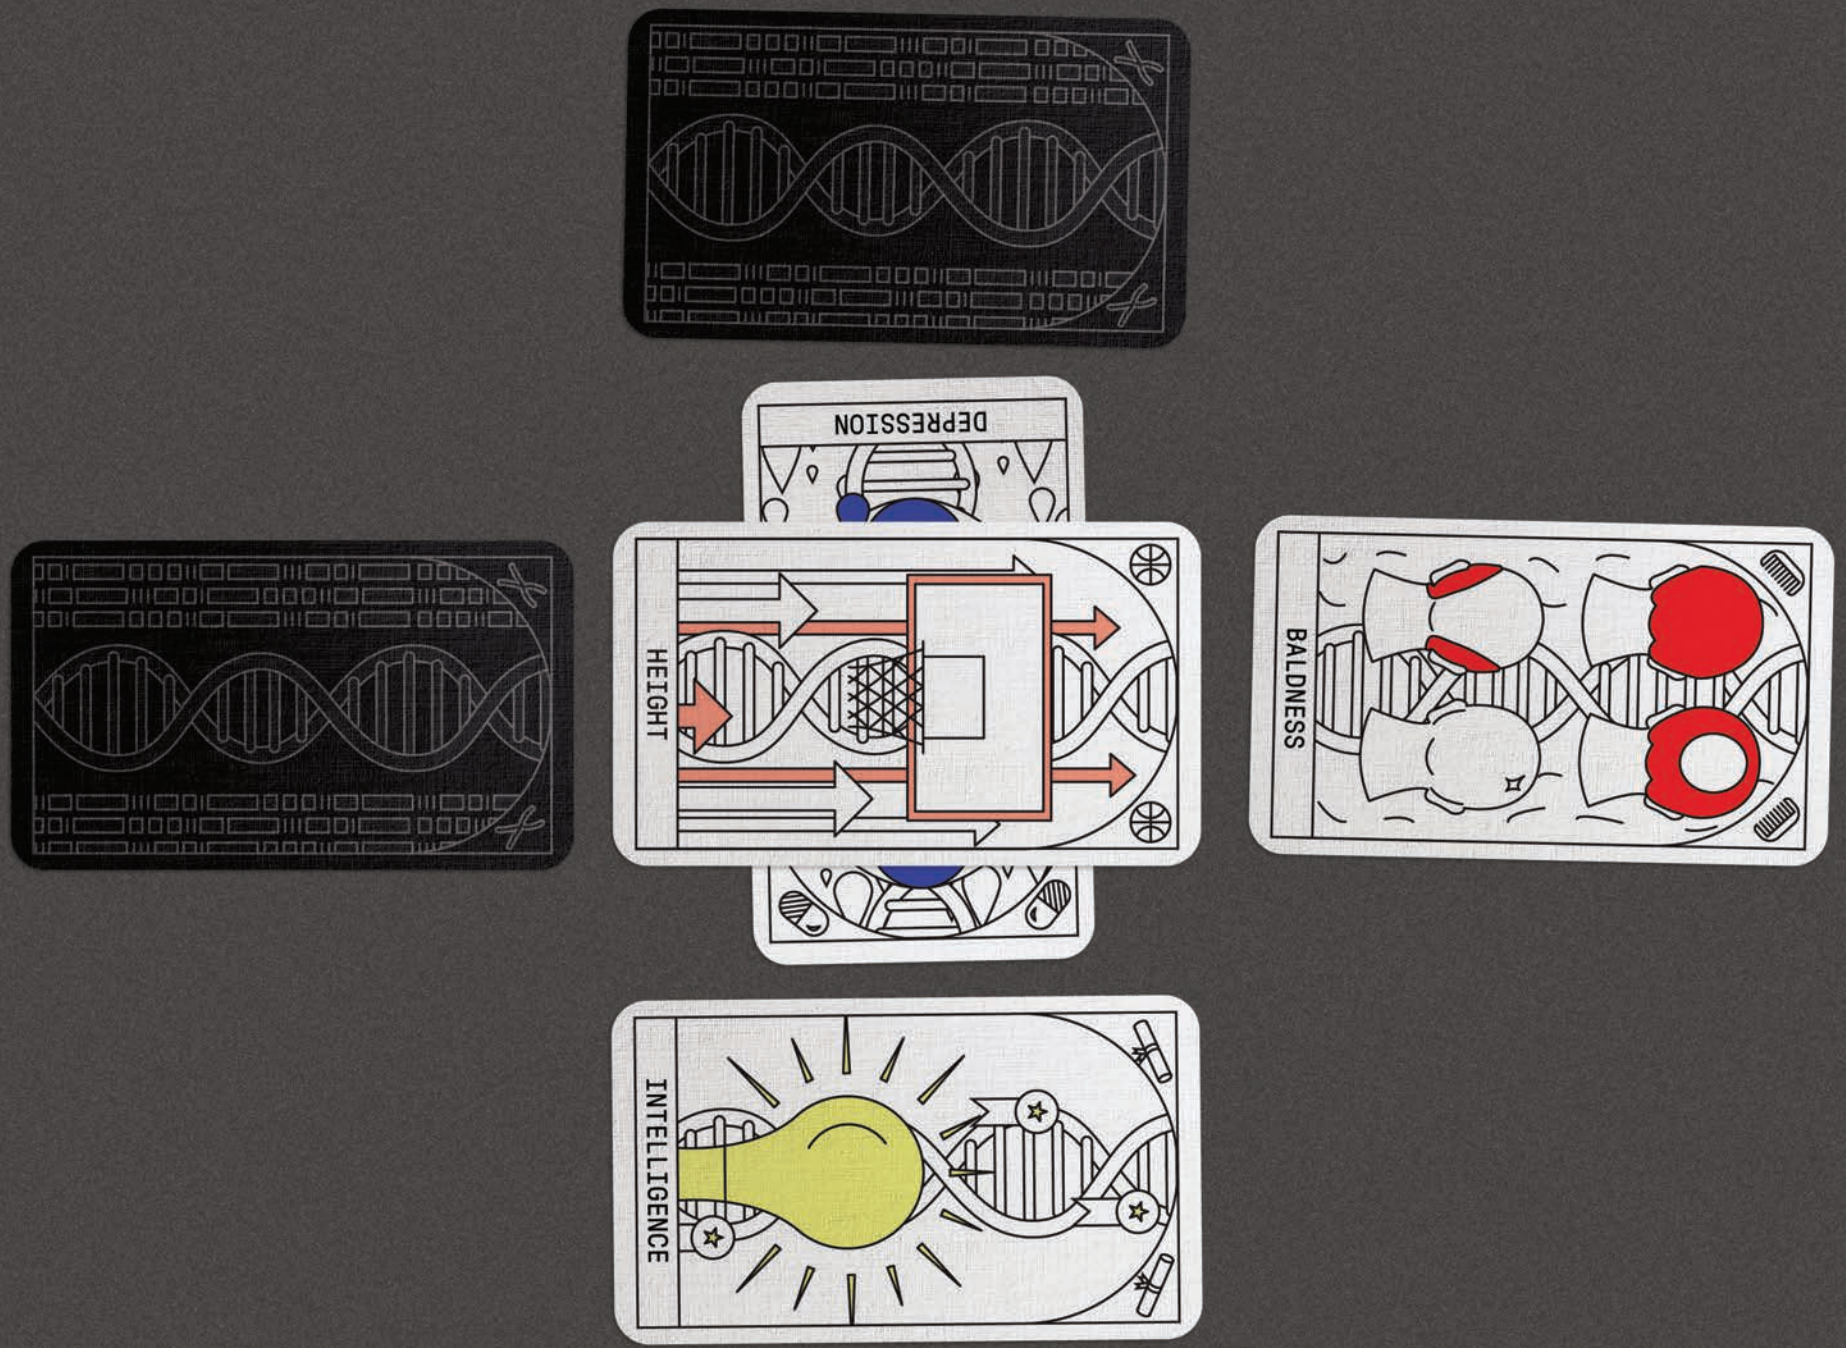

SELMAN DESIGN

Forecasts of genetic fate just got a lot more accurate

DNA-based scores are getting better at predicting intelligence, risks for common diseases, and more.

BY ANTONIO REGALADO

Here's the breakthrough: a new way to guess your chance of serious disease from your DNA. Any drawbacks? You bet. The technology could lead to a society where people get genetic grades at birth.

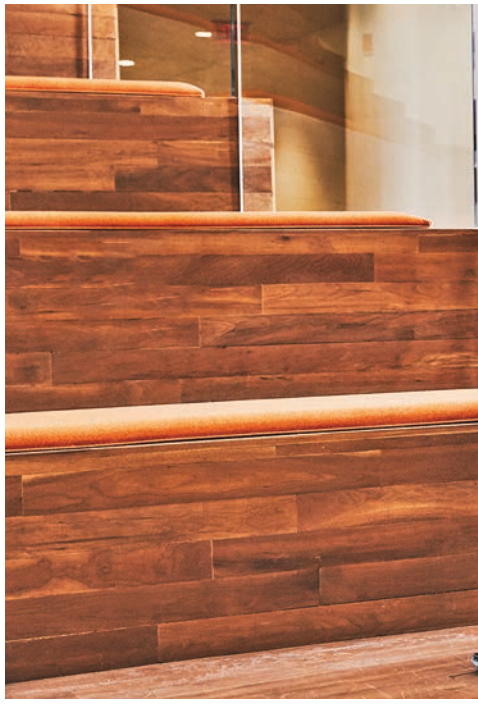

When Amit Khera explains how he predicts disease, the young cardiologist's hands touch the air, arranging imaginary columns of people: 30,000 who have suffered heart attacks here, 100,000 healthy controls there.

There's never been data available on as many people's genes as there is today. And that wealth of information is allowing researchers to guess at any person's chance of getting common diseases like diabetes, arthritis, clogged arteries, and depression.

Doctors already test for rare, deadly mutations in individual genes. Think of the *BRCA* breast cancer gene. Or the one-letter mutation that causes sickle-cell anemia. But such one-to-one connections between a mutation and a disease—"the gene for X"—aren't seen in most common ailments. Instead, these have complex causes, which until recently have remained elusive.

The day I visited him, Khera was constructing what is called a polygenic score—"poly" because his calculations involve thousands of genes, not just one. This particular score predicted a person's chance of developing atrial fibrillation, or irregular heartbeat. It's a common disorder but is often diagnosed only after someone has been rushed to the ER with a stroke.

Khera pointed to his screen. There, seven-digit numbers, each representing an anonymous DNA donor, appeared alongside their scores. The outliers had a risk four times the average.

Khera, who works in the laboratory of heart doctor and gene hunter Sekar Kathiresan at the Broad Institute, in Cambridge, Massachusetts, says the new scores can now identify as much risk for disease as the rare genetic flaws that have preoccupied physicians until now.

"Where I see this going is that at a young age you'll basically get a report card," says Khera. "And it will say for these 10 diseases, here's your score. You are in the 90th percentile for heart disease, 50th for breast cancer, and the lowest 10 percent for diabetes."

Such comprehensive report cards aren't being given out yet, but the science to create them is here. Delving into giant databases like the UK Biobank, which collects the DNA and holds the medical records of some 500,000 Britons, geneticists are peering into the lives of more people and extracting correlations between their genomes and their diseases, personalities, even habits. The latest gene hunt, for the causes of insomnia, involved a record 1,310,010 people.

The sheer quantity of material is what allows scientists like Khera to see how complex patterns of genetic variants are tied to many diseases and traits. Such patterns were hidden in earlier, more limited studies, but now the search for ever smaller signals in ever bigger data is paying off. Give Khera the simplest readout of your genome—the kind created with a $100 DNA-reading chip the size of a theater ticket—and he can add up your vulnerabilities and strengths just as one would a tally in a ledger.

Such predictions, at first hit-or-miss, are becoming more accurate. One test described last year can guess a person's height to within four centimeters, on the basis of 20,000 distinct DNA letters in a genome. As the prediction technology improves, a flood of tests is expected to reach the market. Doctors in California are testing an iPhone app that, if you upload your genetic data, foretells your risk of coronary artery disease. A commercial test launched in September, by Myriad Genetics, estimates the breast cancer chances of any woman, not only the few who have inherited broken versions of the

SIMON SIMARD

Cardiologist Amit Khera is part of a team at the Broad Institute predicting disease from DNA.

BRCA gene. Sharon Briggs, a senior scientist at Helix, which operates an online store for DNA tests, says most of these products will use risk scores within three years.

"It's not that the scores are new," says Briggs. "It's that they're getting much better. There's more data."

TINY INFLUENCES

When they launched the first modern gene searches a decade ago, following the completion of the Human Genome Project, medical researchers still hoped that a few major genetic culprits would explain common diseases like diabetes. "I expect there are about 12 genes involved [in diabetes], and that all of them will be discovered in the next two years," Francis Collins, now the head of the US National Institutes of Health and one of the leading players in sequencing the human genome, confidently declared in 2006.

If that had turned out to be true, the small list of genes would have given drug designers clear, tangible targets. That would easily have justified the whole enterprise, financed with hundreds of millions of US tax dollars. In the case of a few diseases, like macular degeneration, the searches paid off. Mostly, though, geneticists drew in empty nets. By 2009, Collins and others had begun to talk glumly about "the missing heritability."

Where were the disease-causing genes? Everywhere, it turns out. And by 2014, genetic studies were finally big enough to prove it. As the number of people with diabetes who enrolled in the gene search studies rose from 661 to 10,128 to 81,412, the "hits" began rolling in. Instead of 12 genes, we now know, type 2 diabetes is influenced by at least 400 locations in our DNA, and probably many more—each with only a tiny, hard-to-detect effect.

To scientists seeking the ultimate cause of common diseases, that's a huge disappointment. If the causes of diabetes, depression, or schizophrenia are sprinkled around the genome like so much powdered sugar, it means we're far from understanding or curing them. "No one wanted that to be the answer," says Mark Daly, a geneticist at the Broad Institute. "But it is what it is."

While the scattershot nature of inheritance may make disease hard to comprehend, though, the same data is making it much easier to predict. To create their models, Khera and Kathiresan use 6.6 million positions in a person's genome. Each position is a single DNA letter. It could be A for you and G for me. From big genetic studies, Khera can now look up how much more likely a person with a G in that position is to have a heart attack. Maybe it raises the chances by 0.1 percent. That's a negligible amount. Maybe a G in another position reduces the risk by 0.2 percent. But if you add up all the tiny genetic influences, the effect can become substantial.

When they built a predictor for coronary heart disease, for instance, Kathiresan's team discovered that the people it predicted to have the very highest risk, the top 2.5 percent, had four times the average chance of developing clogged arteries. That's about equal to the risk of clogged arteries caused by familial hypercholesterolemia, a condition marked by sky-high cholesterol levels and caused by a single critical gene. If doctors worry about that—which they do—why not also pay attention to the high end of genomic risk scores?

"That's the thing that convinced me," says Kathiresan. And the number of people whose genome predictions raise a red alert will also be much larger. Familial cholesterolemia affects only about one in 250 people. Genome scores would identify about eight times as many people at high risk for heart disease, he believes.

What he's not yet sure about, Kathiresan says, is how to get the new risk information into people's hands. He has considered launching an app or selling the statistical model to a diagnostics company. "Everyone wants to get their score. Everyone is asking where is the product for heart disease," he says. "I tell them, we are working on it."

Heart disease is, in some ways, a best-case scenario for using risk scores. That's because you can change your real-life risk—say, by going on a diet or taking a cholesterol-lowering statin pill. What's more, probabilities are already a big part of heart medicine. Khera, who dons a white coat once a week to treat patients at Massachusetts General Hospital, uses a combination of a person's age, weight, cholesterol levels, and habits like smoking to guess the chance of a heart attack in the next 10 years. Now genetic scores could be added to those models, making them more accurate.

What's powerful about DNA predictions is that they are measurable at any time of life, unlike most risk factors. "If you line up a bunch of 18-year-olds, none of them have high cholesterol, none of them have diabetes. It's a zero in all the columns, and you can't stratify them by who is most at risk," says Khera. "But with a $100 test we can get stratification at least as good as when someone is 50, and for a lot of diseases."

DANGEROUS KNOWLEDGE

Drug companies have started to notice. Last year Anders Dale, a brain researcher at the University of California, San Diego, announced his intention to market a risk calculator for Alzheimer's disease. It will guess whether a person will develop the disease and, if so, at what age.

The service won't launch until this summer, but Dale says drug companies immediately got in touch. Now he is helping three of them test the DNA of people in clinical trials for Alzheimer's drugs (he declined to name them). Despite the billions spent developing such drugs, every one tried so far has flopped. The problem is that when no one knows who will get the disease, it's difficult to know whether a preventive drug is working. If companies could test the drugs only on people with a high risk of Alzheimer's, it would be much easier. It's possible future drugs will be labeled "Recommended for those with polygenic scores 90 and above."

Dale plans to put his Alzheimer's predictor online and charge $99 to anyone who wants to use it. More than ten million people already have their DNA data because they signed up for 23andMe or Ancestry.com to research their family trees. Dale says they will be able to upload the data with a click and receive a report. I asked him why people would want to know ahead of time about a disease that's currently untreatable. "They may want to make plans," he said.

Other doctors believe risk scores will give people the push they need to think harder about their well-being. "I love the idea of polygenic risk scores because the future is health, not medicine," says Steven Tucker, a physician who practices in Singapore. He likes his patients to use wearable devices and trackers, and risk scores could be combined with those. Someone at high risk for atrial fibrillation, for instance, might wear a smart watch with a heart monitor built into it. "My patients want to manage the future," says Tucker. "If you can define it more accurately, there is a better chance you can do something about it."

Even so, the value of the new genomic future-gazing is hotly disputed. That's because the scores are not individual certainties; they are merely rough probabilities derived from large populations. Of people given high scores by Khera's atrial fibrillation predictor, for example, only a small minority, 7 percent, will actually ever develop the condition.

This uncertainty matters because if people are given risk scores, they'll base decisions on them. Last fall, Myriad Genetics became the first large diagnostics company to introduce a polygenic risk test to the US market. Called riskScore, it measures 81 variants to estimate a woman's chance of

DNA-based IQ tests are likely to become available in coming years. Critics fear "some truly dreadful social policies" could result.

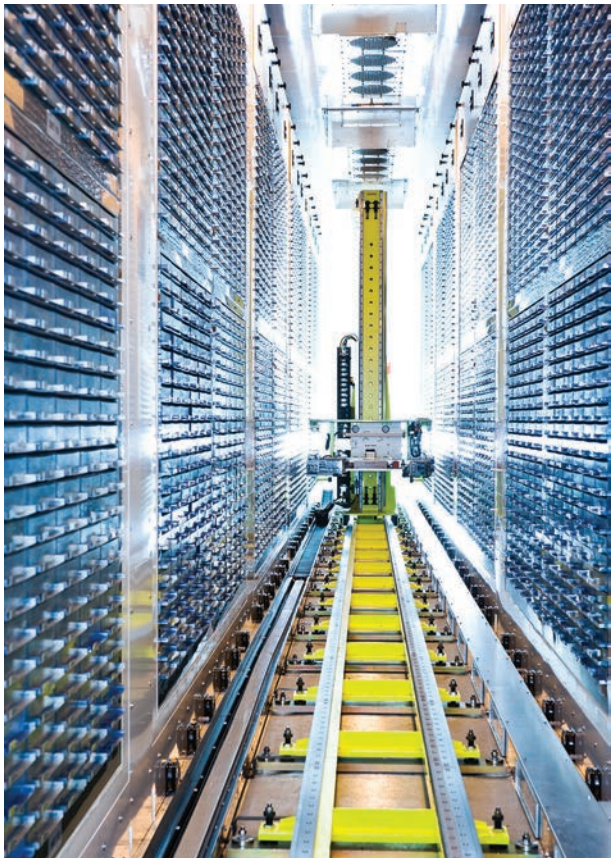

UK BIOBANK

Genetic studies are becoming more powerful as they involve more people. Here, a robot on rails searches among frozen DNA samples donated by 500,000 British volunteers as part of the UK Biobank.

breast cancer. Women with a high score might undergo extra mammograms; those at low risk might skip them. What no one yet knows is whether those decisions will lead to fewer cancer deaths. Finding out will require expensive long-term studies that Myriad, which is selling the test, hasn't yet done.

One physician who finds all this troubling is Patrick Sullivan at the University of North Carolina, Chapel Hill, where he leads the Psychiatric Genomics Consortium. The group has DNA samples from more than 900,000 people with confirmed mental illness, including 60,000 with schizophrenia. This disease is known to be highly influenced by genes. If one identical twin develops schizophrenia, for example, there is a 47 percent chance the other one will too.

But Sullivan says it would be reckless to tell apparently healthy people whether their DNA score for schizophrenia is high or low. Just think of those twins, he says: they have the same DNA and the same score, yet it is even odds a schizophrenia prediction would be wrong. Giving such a flawed forecast to someone "is a terrible idea," says Sullivan. "What you want it to do is distinguish who has it and who doesn't, and we aren't there yet."

A DNA IQ TEST

In addition to predicting disease, geneticists can build models to predict any human trait that can be measured, including behaviors. Is this person destined for a life of crime and recidivism? Will that one be neurotic, depressed, or smarter than average?

The scoring technology, scientists say, will soon shed uncomfortable light on such questions. In January, two leading psychologists argued that direct-to-consumer DNA IQ tests will soon become "routinely available" and will predict children's ability "to learn, reason, and solve problems." They believe parents will test toddlers and use the results to make school plans.

To some, using foggy genetic horoscopes to decide who goes to college and who ends up in trade school sounds like an extraordinarily bad idea. On his blog Gloomy Prospect, Eric Turkheimer, a prominent psychologist at the University of Virginia, says the danger is that the scores will be overinterpreted to "recommend some truly dreadful social policies." That, he thinks, would be "the worst possible kind of biologically determinist discrimination." To Turkheimer, polygenic scores are "less than meets the eye" and about as fair as "predicting your IQ from a cousin you haven't met."

Such views aren't stopping the rapid pace of genetic exploration. Until last year, no gene variant had ever been tied directly to IQ test results. Since then, studies involving more than 300,000 people's DNA have linked 206 variants to intelligence. It means genetic scores can now account for 10 percent of a person's performance on an IQ test. That could reach 25 percent within a few years, as more data accumulates. One US company, Genomic Prediction, even says it wants to test IVF embryos for intelligence, so parents can discard those expected to be mentally unfit.

Dystopia, dubious medicine, or a breakthrough in prevention? Genomic prediction may well be all three. What is clear is that, with the data needed to create predictors becoming freely available online, 2018 will be a breakout year for DNA fortune-telling.

Antonio Regalado is MIT Technology Review*'s senior editor for biomedicine.*

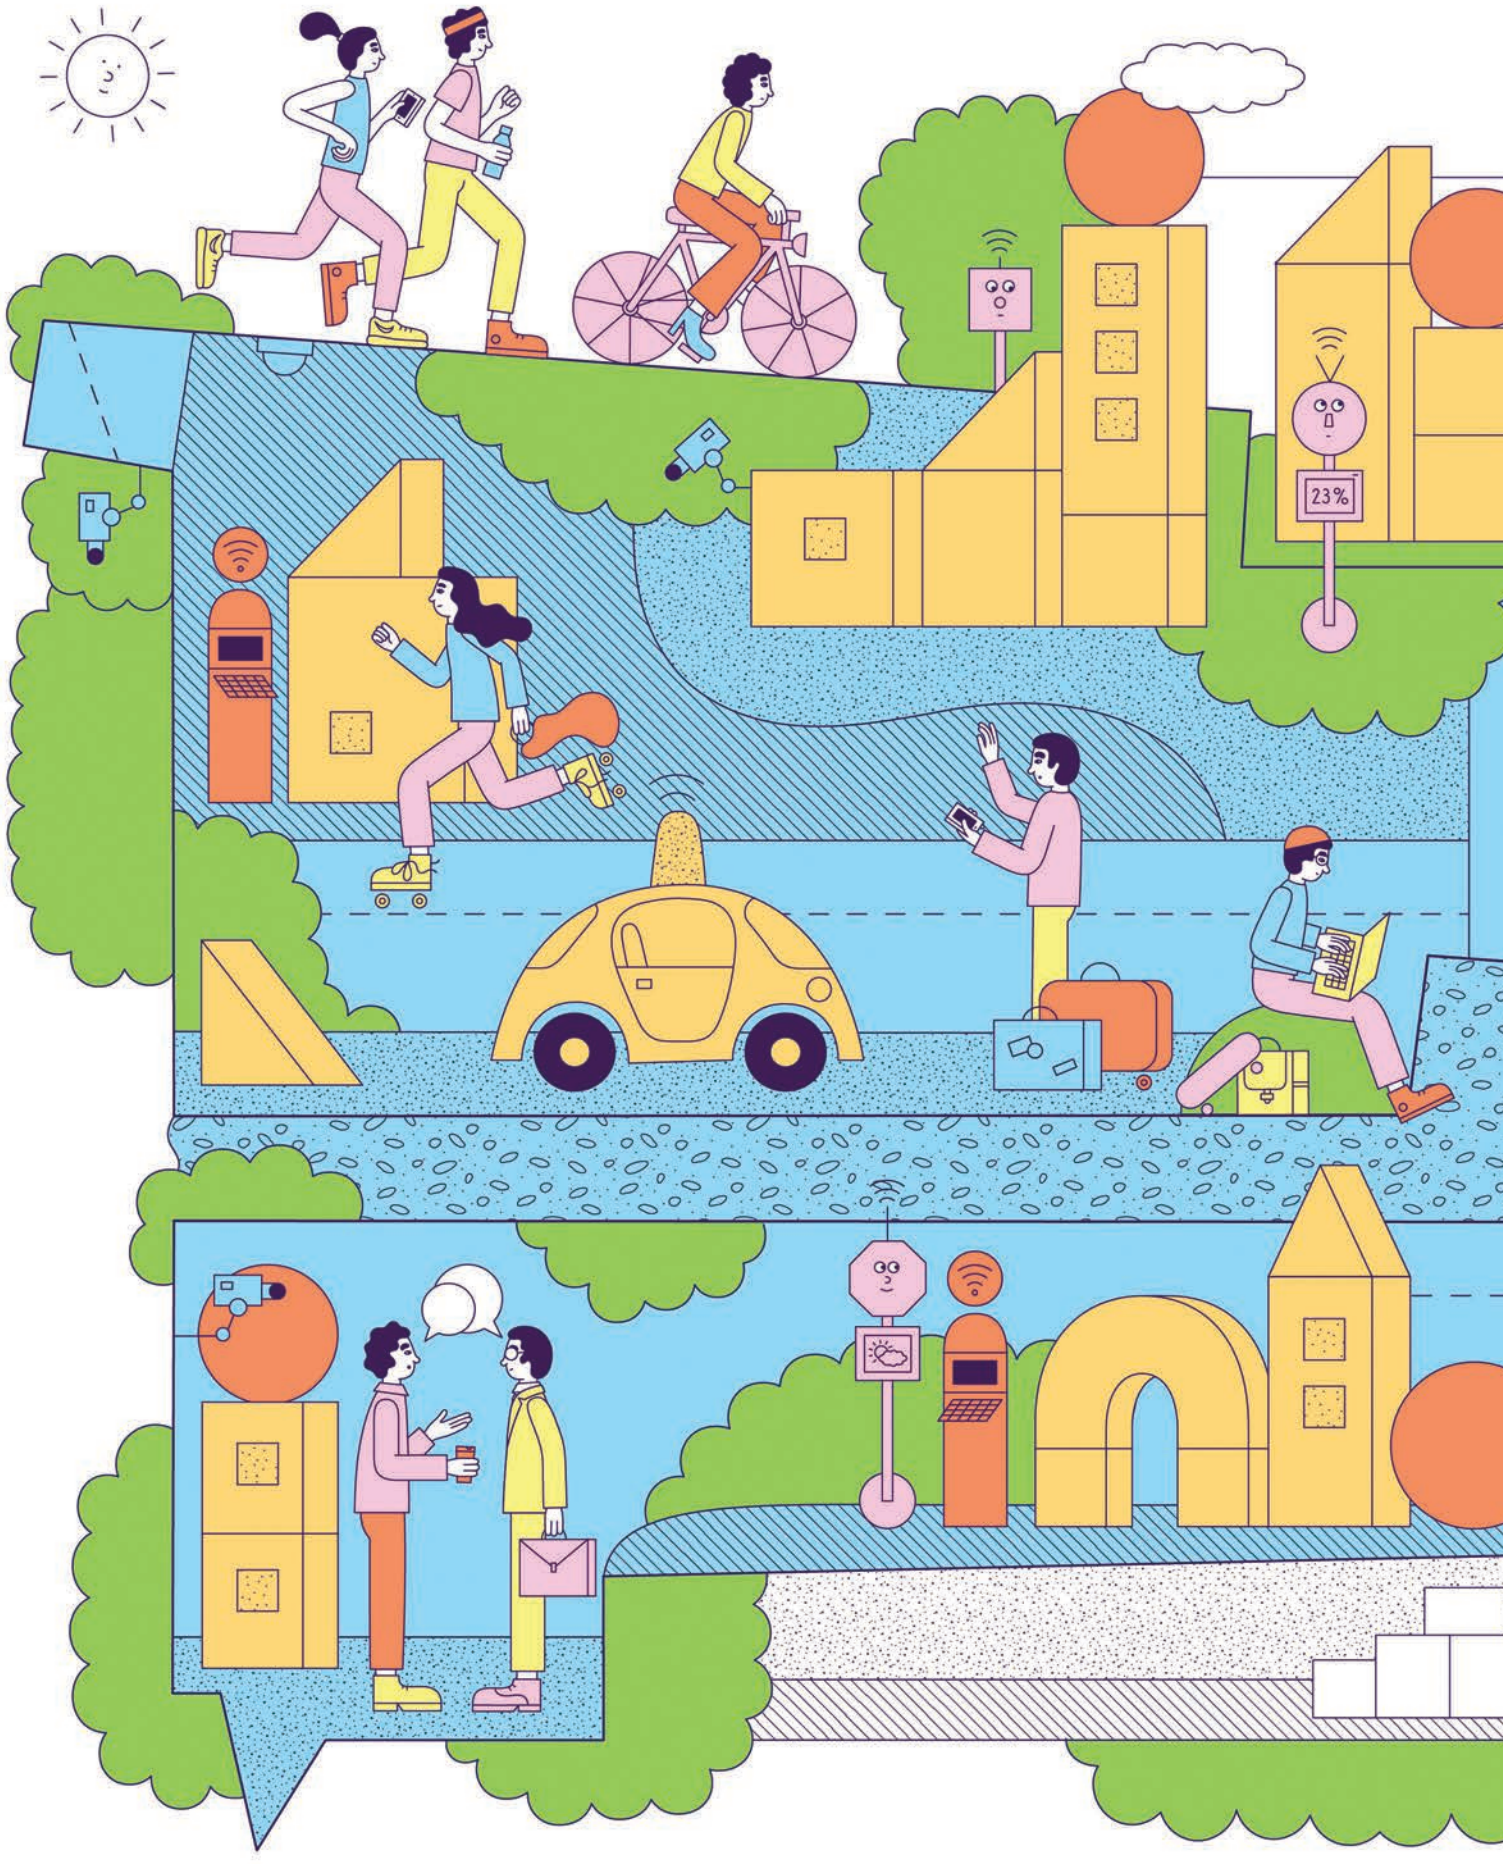
23%

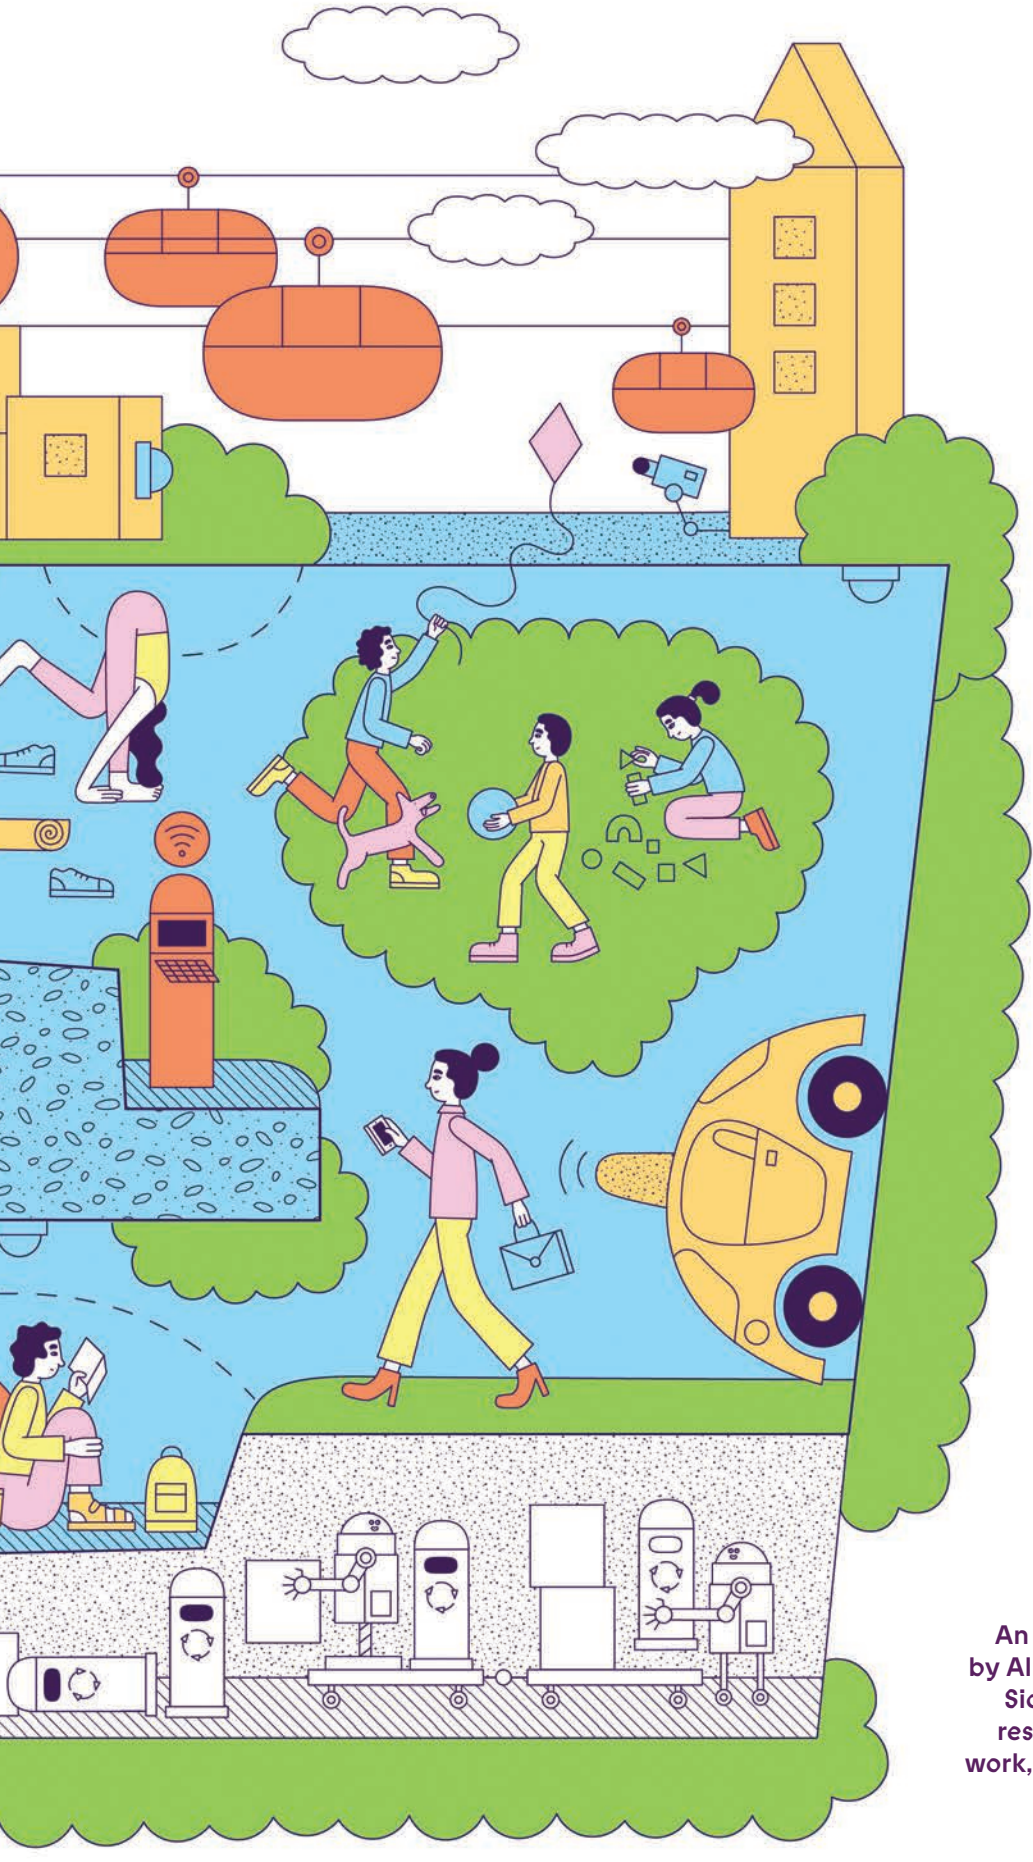

A SMARTER SMART CITY

By Elizabeth Woyke

Illustrations by Martina Paukova

An ambitious project by Alphabet subsidiary Sidewalk Labs could reshape how we live, work, and play in urban neighborhoods.

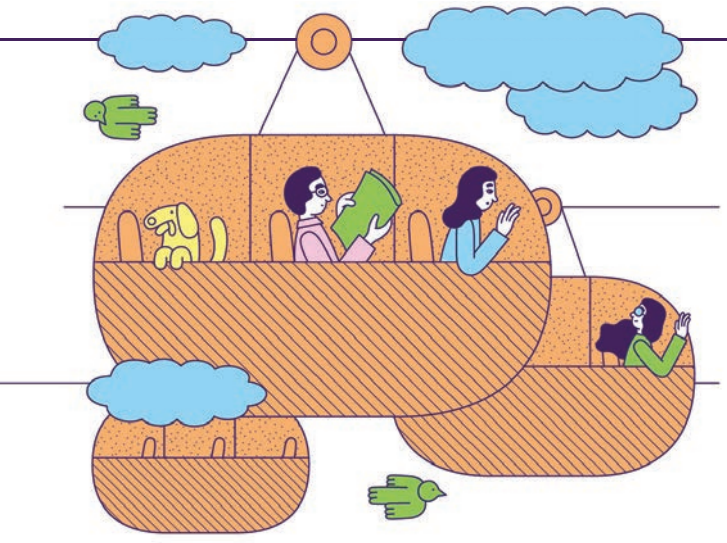

On Toronto's waterfront, where the eastern part of the city meets Lake Ontario, is a patchwork of cement and dirt. It's home to plumbing and electrical supply shops, parking lots, winter boat storage, and a hulking silo built in 1943 to store soybeans—a relic of the area's history as a shipping port.

Torontonians describe the site as blighted, underutilized, and contaminated. Alphabet's Sidewalk Labs wants to transform it into one of the world's most innovative city neighborhoods. It will, in the company's vision, be a place where driverless shuttle buses replace private cars; traffic lights track the flow of pedestrians, bicyclists, and vehicles; robots transport mail and garbage via underground tunnels; and modular buildings can be expanded to accommodate growing companies and families.

In the early 2000s, so-called smart cities were all the rage. Captivated by the idea of urban districts that would use technology to reduce energy consumption and pollution, make transportation more efficient, and lure affluent tenants, countries including China, South Korea, and the United Arab Emirates hired developers to transform large swaths of land into photogenic cities stuffed with the latest innovations.

All fell short of their lofty ambitions. Sidewalk Labs, which was founded in 2015 as a subsidiary of Alphabet to develop technology for alleviating urban problems, believes it can buck the trend by working closely with the community and tailoring the technology to local needs. "People have been trying to build the city of the future for more than 100 years," says Rit Aggarwala, the executive in charge of Sidewalk Labs' urban-systems planning. "But we really want to tap into [Toronto's] existing vitality and character."

The neighborhood, called Quayside, is Sidewalk's first big project. It will start life on a 12-acre plot mostly owned by Waterfront Toronto—a local development agency founded by Canada's federal, provincial, and municipal governments—and is expected to house an estimated 5,000 people. Later development could expand to a neighboring 700-plus-acre parcel of industrial waterfront and involve tens of thousands of residents. "All of our thinking and decisions on Quayside are shaped by the question 'What do 21st-century technologies enable us to do better?'" says Aggarwala.

Driverless cars will have a big role. Sidewalk assumes they will navigate more precisely and obey traffic laws more consistently than human drivers, so it wants to put narrower lanes in Quayside and carve out more room for sidewalks and parks. In theory, using shared self-driving vehicles will mean that fewer people need to own cars, saving families a projected $6,000 a year.

Sensing and monitoring public activity accurately and frequently will be key. Running autonomous buses on city streets requires knowing when to change lights and other signals to give cyclists and pedestrians priority.

Sidewalk Labs says the sensor information would also support long-term planning. The data would fuel a virtual model of Quayside, which urban planners could use to test infrastructure changes quickly, at low cost, and without bothering residents. It

"All of our thinking and decisions on Quayside are shaped by the question 'What do 21st-century technologies enable us to do better?'"

An aerial view shows the land to be included in Quayside, the tech-centric neighborhood Sidewalk Labs and Waterfront Toronto plan to build.

SIDEWALK LABS

could also be stored in a shared repository that entrepreneurs and companies could draw on to make their own products and services for Quayside.

Waterfront Toronto issued a request for a partner in March 2017 and announced its choice of Sidewalk Labs in October. The two organizations now have a year to decide what technologies to deploy, which companies besides Sidewalk will provide them, and how to finance the project. Sidewalk has pledged to spend $50 million on the first phase of planning, which has already kicked off, and pilot testing, which is scheduled to begin later this year. Construction could start next year.

"If Quayside is successful, cities in Canada and elsewhere won't have any choice but to emulate it, because it would have enormous implications in terms of sustainability and quality of life," says Christopher De Sousa, who directs the school of urban and regional planning at Ryerson University in Toronto.

CITY AS SMARTPHONE

Unsurprisingly for a company spawned, in part, by technologists, Sidewalk thinks of smart cities as being rather like smartphones. It sees itself as a platform provider responsible for offering basic tools (from software that identifies available parking spots to location-based services monitoring the exact position of delivery robots), much as Google does with its smartphone operating system, Android. Details are still under discussion, but Sidewalk plans to let third parties access the data and technologies, just as developers can use Google's and Apple's software tools to craft apps.

In fact, Sidewalk anticipates that 80 percent of the work on Quayside will involve these third parties. Some are likely to be other Alphabet companies, such as the autonomous-vehicle maker Waymo, but Sidewalk has said that competitors such as the ride-hailing provider Lyft would be able to operate in Quayside, too.

That will help Sidewalk tailor its products to cities around the world. "If you think of the city as a platform and design in the ability for people to change it as quickly as you and I can customize our iPhones, you make it authentic because it doesn't just reflect a central plan," Aggarwala says. "It also reflects the people who live and work there."

UBIQUITOUS SENSORS

One way Sidewalk Labs differs from Google and Android is the way it plans to make money. The startup is still figuring out its business model, but it has said it expects to license the technology it produces for Quayside to other cities, rather than compile consumer data and use it to sell ads. Still, many people are wor-

"If Quayside is successful, cities in Canada and elsewhere won't have any choice but to emulate it."

ried about privacy given the data Quayside plans to harvest. In recent months, Canadian newspapers and blogs have published a number of skeptical articles, including a list of more than 35 probing questions about Sidewalk Labs' business model and data governance methods.

Most of the scrutiny relates to the company's plan to install sensors across Quayside to measure everything from building occupancy to sewage flow rates to how often a public waste bin is used. The company is also devising a system that would give Quayside residents and workers a quick way to pay for things and gain access to services, similar to the way people can buy products off Amazon with a single click or get billed for Uber rides through an app.

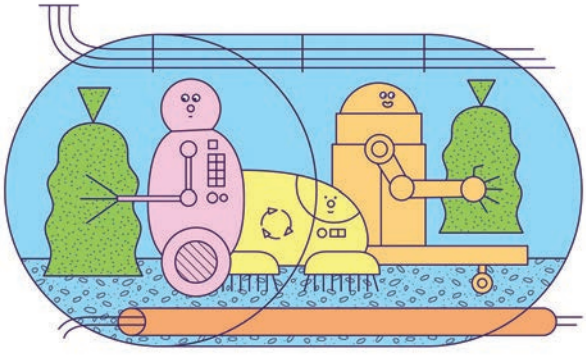

TRASH ROBOTS
Sidewalk Labs proposes transporting garbage in underground tunnels to reduce street traffic and greenhouse-gas emissions. Robots would sort and haul the waste and could also deliver mail and packages.

Though Sidewalk Labs says the data would be used for a community purpose, such as giving transit discounts to low-income residents, regulating building temperatures, and keeping trash cans from overflowing, not everyone is convinced. "There are definitely questions about whether Sidewalk Labs will try to make money by tracking people's daily interactions," says David Roberts, who studies cities at the University of Toronto. "What data will be collected, how personal will it be, how will it be used, and who will have access to it?"

Aggarwala says Sidewalk will gather only the data needed to solve the problem at hand. If the company wanted to analyze pedestrian patterns in Quayside, it could employ a lidar device, which uses light from a laser to detect objects; or a camera with very low resolution; or one that counts people who walk by but doesn't store the images. "That way you can capture the information you need but not anything that can be traced back to an individual," he says. "I think if we can do that and demonstrate how that can improve urban life, then people will be comfortable and give us the license to do it."

Waterfront Toronto says it will hold Sidewalk accountable for explaining to the public what personal data it collects and why, and how it keeps the information secure. It will require that the data be stored in Canada. The agency also says that Sidewalk Labs won't share Quayside data with Google automatically, even though Google may contribute technology to the project through its maps or cloud computing services.

To win residents over, Sidewalk Labs and Waterfront Toronto held a town hall meeting in November and will have neighborhood discussions, expert roundtables, and design competitions in the coming months. The project will succeed only if it convinces Torontonians that it offers solutions to urban problems, especially since the city's downtown population is expected to almost double, to nearly half a million people, over the next 20 years.

"Quayside can't be a rich person's enclave or a detached tech corridor," says Will Fleissig, the head of Waterfront Toronto. "It should be a place where anyone would want to work, hang out, and raise a family."

Sidewalk Labs clearly has some skepticism to overcome. But if it can demonstrate that intense data-gathering really does enhance urban living, it could provide a model for smart cities around the world. Says Matti Siemiatycki, an expert on geography and planning at the University of Toronto, "With these folks, what's intriguing is the sky's the limit in terms of their imagination, the scale they dream at, and the resources they have at their fingertips."

Elizabeth Woyke is MIT Technology Review*'s senior business editor.*

Hello, Quantum

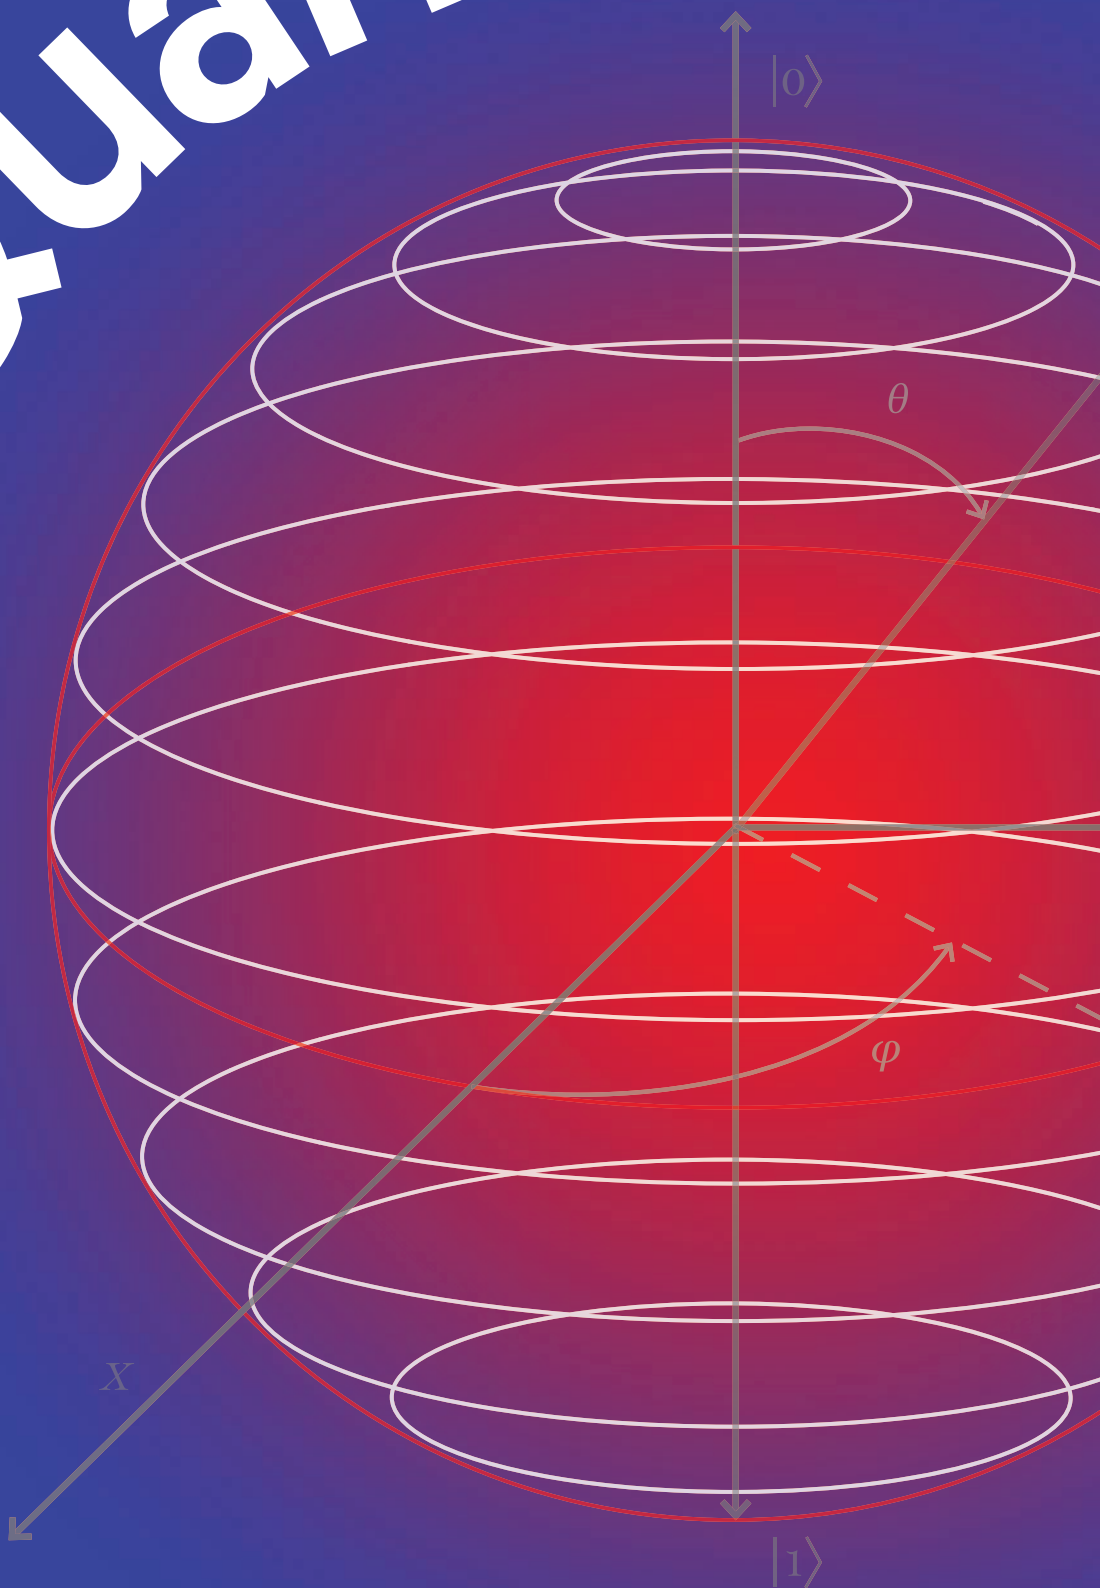

World

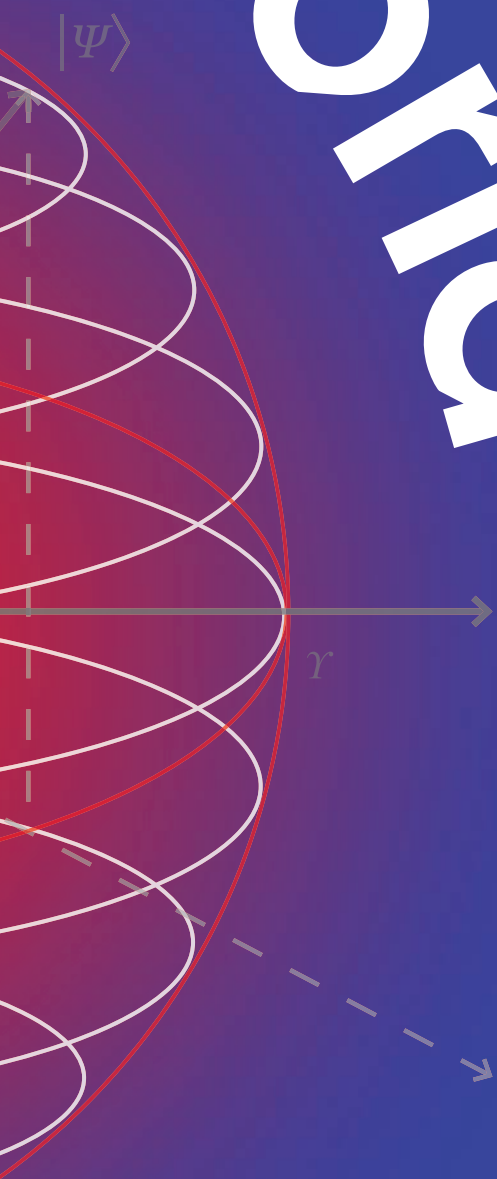

Quantum computers are finally here. What are we going to do with them?

by Will Knight

Inside a small laboratory in lush countryside about 50 miles north of New York City, an elaborate tangle of tubes and electronics dangles from the ceiling. This mess of equipment is a computer. Not just any computer, but one on the verge of passing what may, perhaps, go down as one of the most important milestones in the history of the field.

Quantum computers promise to run calculations far beyond the reach of any conventional supercomputer. They might revolutionize the discovery of new materials by making it possible to simulate the behavior of matter down to the atomic level. Or they could upend cryptography and security by cracking otherwise invincible codes. There is even hope they will supercharge artificial intelligence

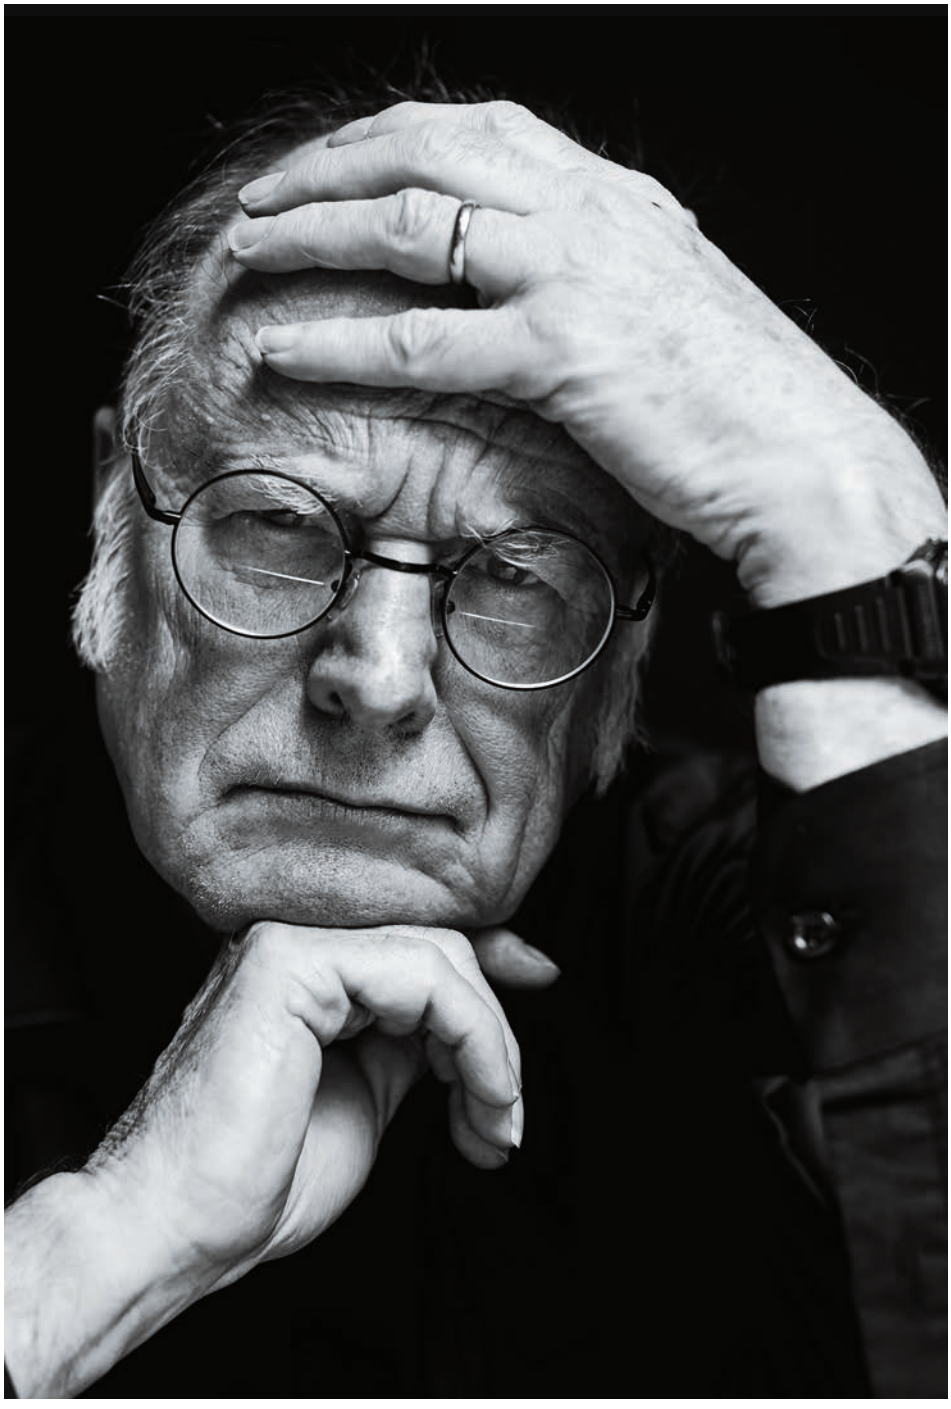

Charles Bennett of IBM Research is one of the founding fathers of quantum information theory. His work at IBM helped create a theoretical foundation for quantum computing.

BARTEK SADOWSKI

by crunching through data more efficiently.

Yet only now, after decades of gradual progress, are researchers finally close to building quantum computers powerful enough to do things that conventional computers cannot. It's a landmark somewhat theatrically dubbed "quantum supremacy." Google has been leading the charge toward this milestone, while Intel and Microsoft also have significant quantum efforts. And then there are well-funded startups including Rigetti Computing, IonQ, and Quantum Circuits.

No other contender can match IBM's pedigree in this area, though. Starting 50 years ago, the company produced advances in materials science that laid the foundations for the computer revolution. Which is why, last October, I found myself at IBM's Thomas J. Watson Research Center to try to answer these questions: What, if anything, will a quantum computer be good for? And can a practical, reliable one even be built?

Why we think we need a quantum computer

The research center, located in Yorktown Heights, looks a bit like a flying saucer as imagined in 1961. It was designed by the neo-futurist architect Eero Saarinen and built during IBM's heyday as a maker of large mainframe business machines. IBM was the world's largest computer company, and within a decade of the research center's construction it had become the world's fifth-largest company of any kind, just behind Ford and General Electric.

While the hallways of the building look out onto the countryside,

the design is such that none of the offices inside have any windows. It was in one of these cloistered rooms that I met Charles Bennett. Now in his 70s, he has large white sideburns, wears black socks with sandals, and even sports a pocket protector with pens in it. Surrounded by old computer monitors, chemistry models, and, curiously, a small disco ball, he recalled the birth of quantum computing as if it were yesterday.

When Bennett joined IBM in 1972, quantum physics was already half a century old, but computing still relied on classical physics and the mathematical theory of information that Claude Shannon had developed at MIT in the 1950s. It was Shannon who defined the quantity of information in terms of the number of "bits" (a term he popularized but did not coin) required to store it. Those bits, the *0*s and *1*s of binary code, are the basis of all conventional computing.

A year after arriving at Yorktown Heights, Bennett helped lay the foundation for a quantum information theory that would challenge all that. It relies on exploiting the peculiar behavior of objects at the atomic scale. At that size, a particle can exist "superposed" in many states (e.g., many different positions) at once. Two particles can also exhibit "entanglement," so that changing the state of one may instantaneously affect the other.

Bennett and others realized that some kinds of computations that are exponentially time consuming, or even impossible, could be efficiently performed with the help of quantum phenomena. A quantum computer would store information in quantum bits, or qubits. Qubits can exist in superpositions of *1* and *0*, and entanglement and a trick called interference can be used to find the solution to a computation over an exponentially large number of states. It's annoyingly hard to compare quantum and classical computers, but roughly speaking, a quantum computer with just a few hundred qubits would be able to perform more calculations simultaneously than there are atoms in the known universe.

"Nature is quantum, goddamn it! So if we want to simulate it, we need a quantum computer."

In the summer of 1981, IBM and MIT organized a landmark event called the First Conference on the Physics of Computation. It took place at Endicott House, a French-style mansion not far from the MIT campus.

In a photo that Bennett took during the conference, several of the most influential figures from the history of computing and quantum physics can be seen on the lawn, including Konrad Zuse, who developed the first programmable computer, and Richard Feynman, an important contributor to quantum theory. Feynman gave the conference's keynote speech, in which he raised the idea of computing using quantum effects. "The biggest boost quantum information theory got was from Feynman," Bennett told me. "He said, 'Nature is quantum, goddamn it! So if we want to simulate it, we need a quantum computer.'"

IBM's quantum computer—one of the most promising in existence—is located just down the hall from Bennett's office. The machine is designed to create and manipulate the essential element in a quantum computer: the qubits that store information.

The gap between the dream and the reality

The IBM machine exploits quantum phenomena that occur in superconducting silicon. Within superconducting silicon, for instance, current will flow clockwise and counterclockwise at the same time. IBM's computer uses superconducting circuits in which two distinct electromagnetic energy states make up a qubit.

The superconducting approach has key advantages. The hardware can be made using well-established manufacturing methods, and a conventional computer can be used to control the system. The qubits in a superconducting circuit are also easier to manipulate and less delicate than individual photons or ions.

Inside IBM's quantum lab, engineers are working on a version of the computer with 50 qubits. You can run a simulation of a simple quantum computer on a normal computer, but at around 50 qubits it becomes nearly impossible. That means IBM is theoretically approaching the point where a quantum computer can solve problems a classical computer cannot: in other words, quantum supremacy.

But as IBM's researchers will tell you, quantum supremacy is an elusive concept. You would need all 50 qubits to work perfectly, when in reality quantum computers are beset by errors that need to be corrected for. It is also devilishly difficult to maintain qubits for any length of time; they tend to "decohere," or lose their delicate quantum nature, much as a smoke ring breaks up at the slightest air current. And the more qubits, the harder both challenges become.

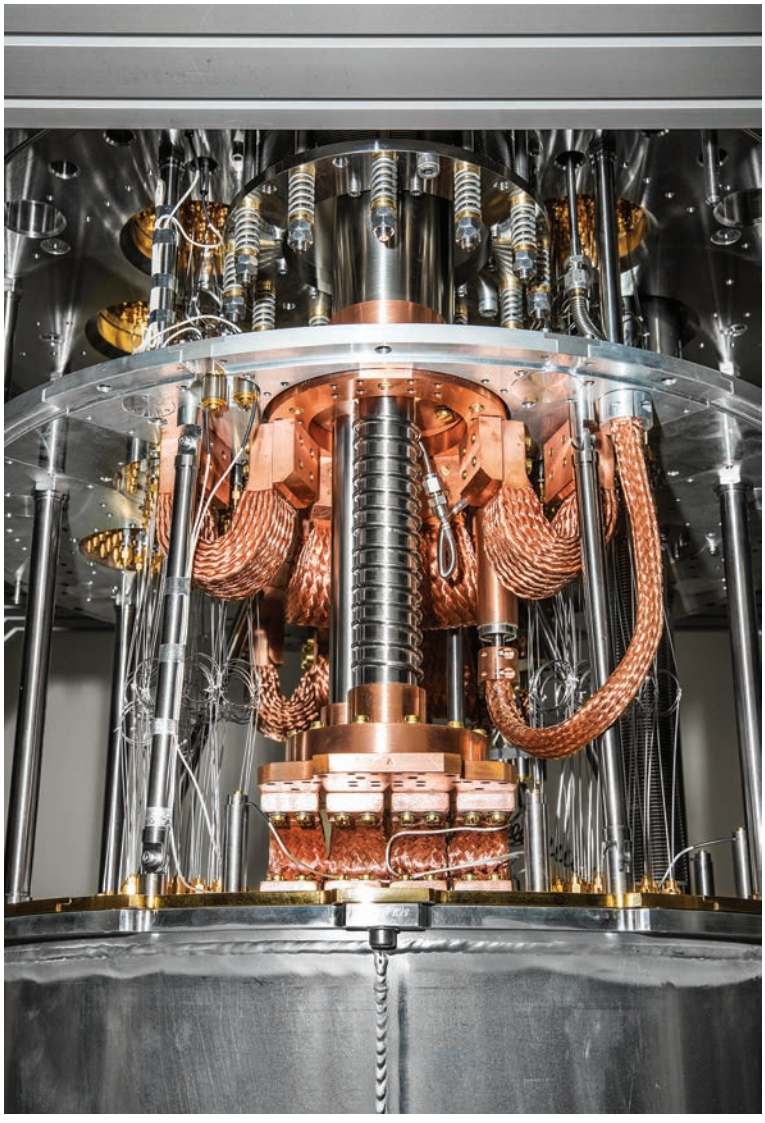

The chips inside IBM's quantum computer (at bottom) are cooled to 15 millikelvin.

"If you had 50 or 100 qubits and they really worked well enough, and were fully error-corrected—you could do unfathomable calculations that can't be replicated on any classical machine, now or ever," says Robert Schoelkopf, a Yale professor and founder of a company called Quantum Circuits. "The flip side to quantum computing is that there are exponential ways for it to go wrong."

Another reason for caution is that it isn't obvious how useful even a perfectly functioning quantum computer would be. It doesn't simply speed up any task you throw at it; in fact, for many calculations, it would actually be slower than classical machines. Only a handful of algorithms have so far been devised where a quantum computer would clearly have an edge. And even for those, that edge might be short-lived. The most famous quantum algorithm, developed by Peter Shor at MIT, is for finding the prime factors of an integer. Many common cryptographic schemes rely on the fact that this is hard for a conventional computer to do. But cryptography could adapt, creating new kinds of codes that don't rely on factorization.

This is why, even as they near the 50-qubit milestone, IBM's own researchers are keen to dispel the hype around it. At a table in the hallway that looks out onto the lush lawn outside, I encountered Jay Gambetta, a tall, easy-going Australian who researches quantum algorithms and potential applications for IBM's hardware. "We're at this unique stage," he said, choosing his words with care. "We have this device that is more complicated than you can simulate on a classical computer, but it's not yet controllable to the precision that you could do the algorithms you know how to do."

What gives the IBMers hope is that even an imperfect quantum computer might still be a useful one.

Gambetta and other researchers have zeroed in on an application that Feynman envisioned back in 1981. Chemical reactions and the properties of materials are determined by the interactions between atoms and molecules. Those interactions are governed by quantum phenomena. A quantum computer can—at least in theory—model those in a way a conventional one cannot.

Last year, Gambetta and colleagues at IBM used a seven-qubit machine to simulate the precise structure of beryllium hydride. At just three atoms, it is the most com-

"The thing driving the hype is the realization that quantum computing is actually real. It is no longer a physicist's dream—it is an engineer's nightmare."

JEREMY LIEBMAN

plex molecule ever modeled with a quantum system. Ultimately, researchers might use quantum computers to design more efficient solar cells, more effective drugs, or catalysts that turn sunlight into clean fuels.

Those goals are a long way off. But, Gambetta says, it may be possible to get valuable results from an error-prone quantum machine paired with a classical computer.

From a physicist's dream to an engineer's nightmare

"The thing driving the hype is the realization that quantum computing is actually real," says Isaac Chuang, a lean, soft-spoken MIT professor. "It is no longer a physicist's dream—it is an engineer's nightmare."

Chuang led the development of some of the earliest quantum computers, working at IBM in Almaden, California, during the late 1990s and early 2000s. Though he is no longer working on them, he thinks we are at the beginning of something very big—that quantum computing will eventually even play a role in artificial intelligence.

But he also suspects that the revolution will not really begin until a new generation of students and hackers get to play with practical machines. Quantum computers require not just different programming languages but a fundamentally different way of thinking about what programming is. As Gambetta puts it: "We don't really know what the equivalent of 'Hello, world' is on a quantum computer."

We are beginning to find out. In 2016 IBM connected a small quantum computer to the cloud. Using a programming tool kit called QISket, you can run simple programs on it; thousands of people, from academic researchers to schoolkids, have built QISket programs that run basic quantum algorithms. Now Google and other companies are also putting their nascent quantum computers online. You can't do much with them, but at least they give people outside the leading labs a taste of what may be coming.

The startup community is also getting excited. A short while after seeing IBM's quantum computer, I went to the University of Toronto's business school to sit in on a pitch competition for quantum startups. Teams of entrepreneurs nervously got up and presented their ideas to a group of professors and investors. One company hoped to use quantum computers to model the financial markets. Another planned to have them design new proteins. Yet another wanted to build more advanced AI systems. What went unacknowledged in the room was that each team was proposing a business built on a technology so revolutionary that it barely exists. Few seemed daunted by that fact.

This enthusiasm could sour if the first quantum computers are slow to find a practical use. The best guess from those who truly know the difficulties—people like Bennett and Chuang—is that the first useful machines are still several years away. And that's assuming the problem of managing and manipulating a large collection of qubits won't ultimately prove intractable.

Still, the experts hold out hope. When I asked him what the world might be like when my two-year-old son grows up, Chuang, who learned to use computers by playing with microchips, responded with a grin. "Maybe your kid will have a kit for building a quantum computer," he said.

Will Knight is senior editor for AI.

This lab at IBM houses quantum machines connected to the cloud.

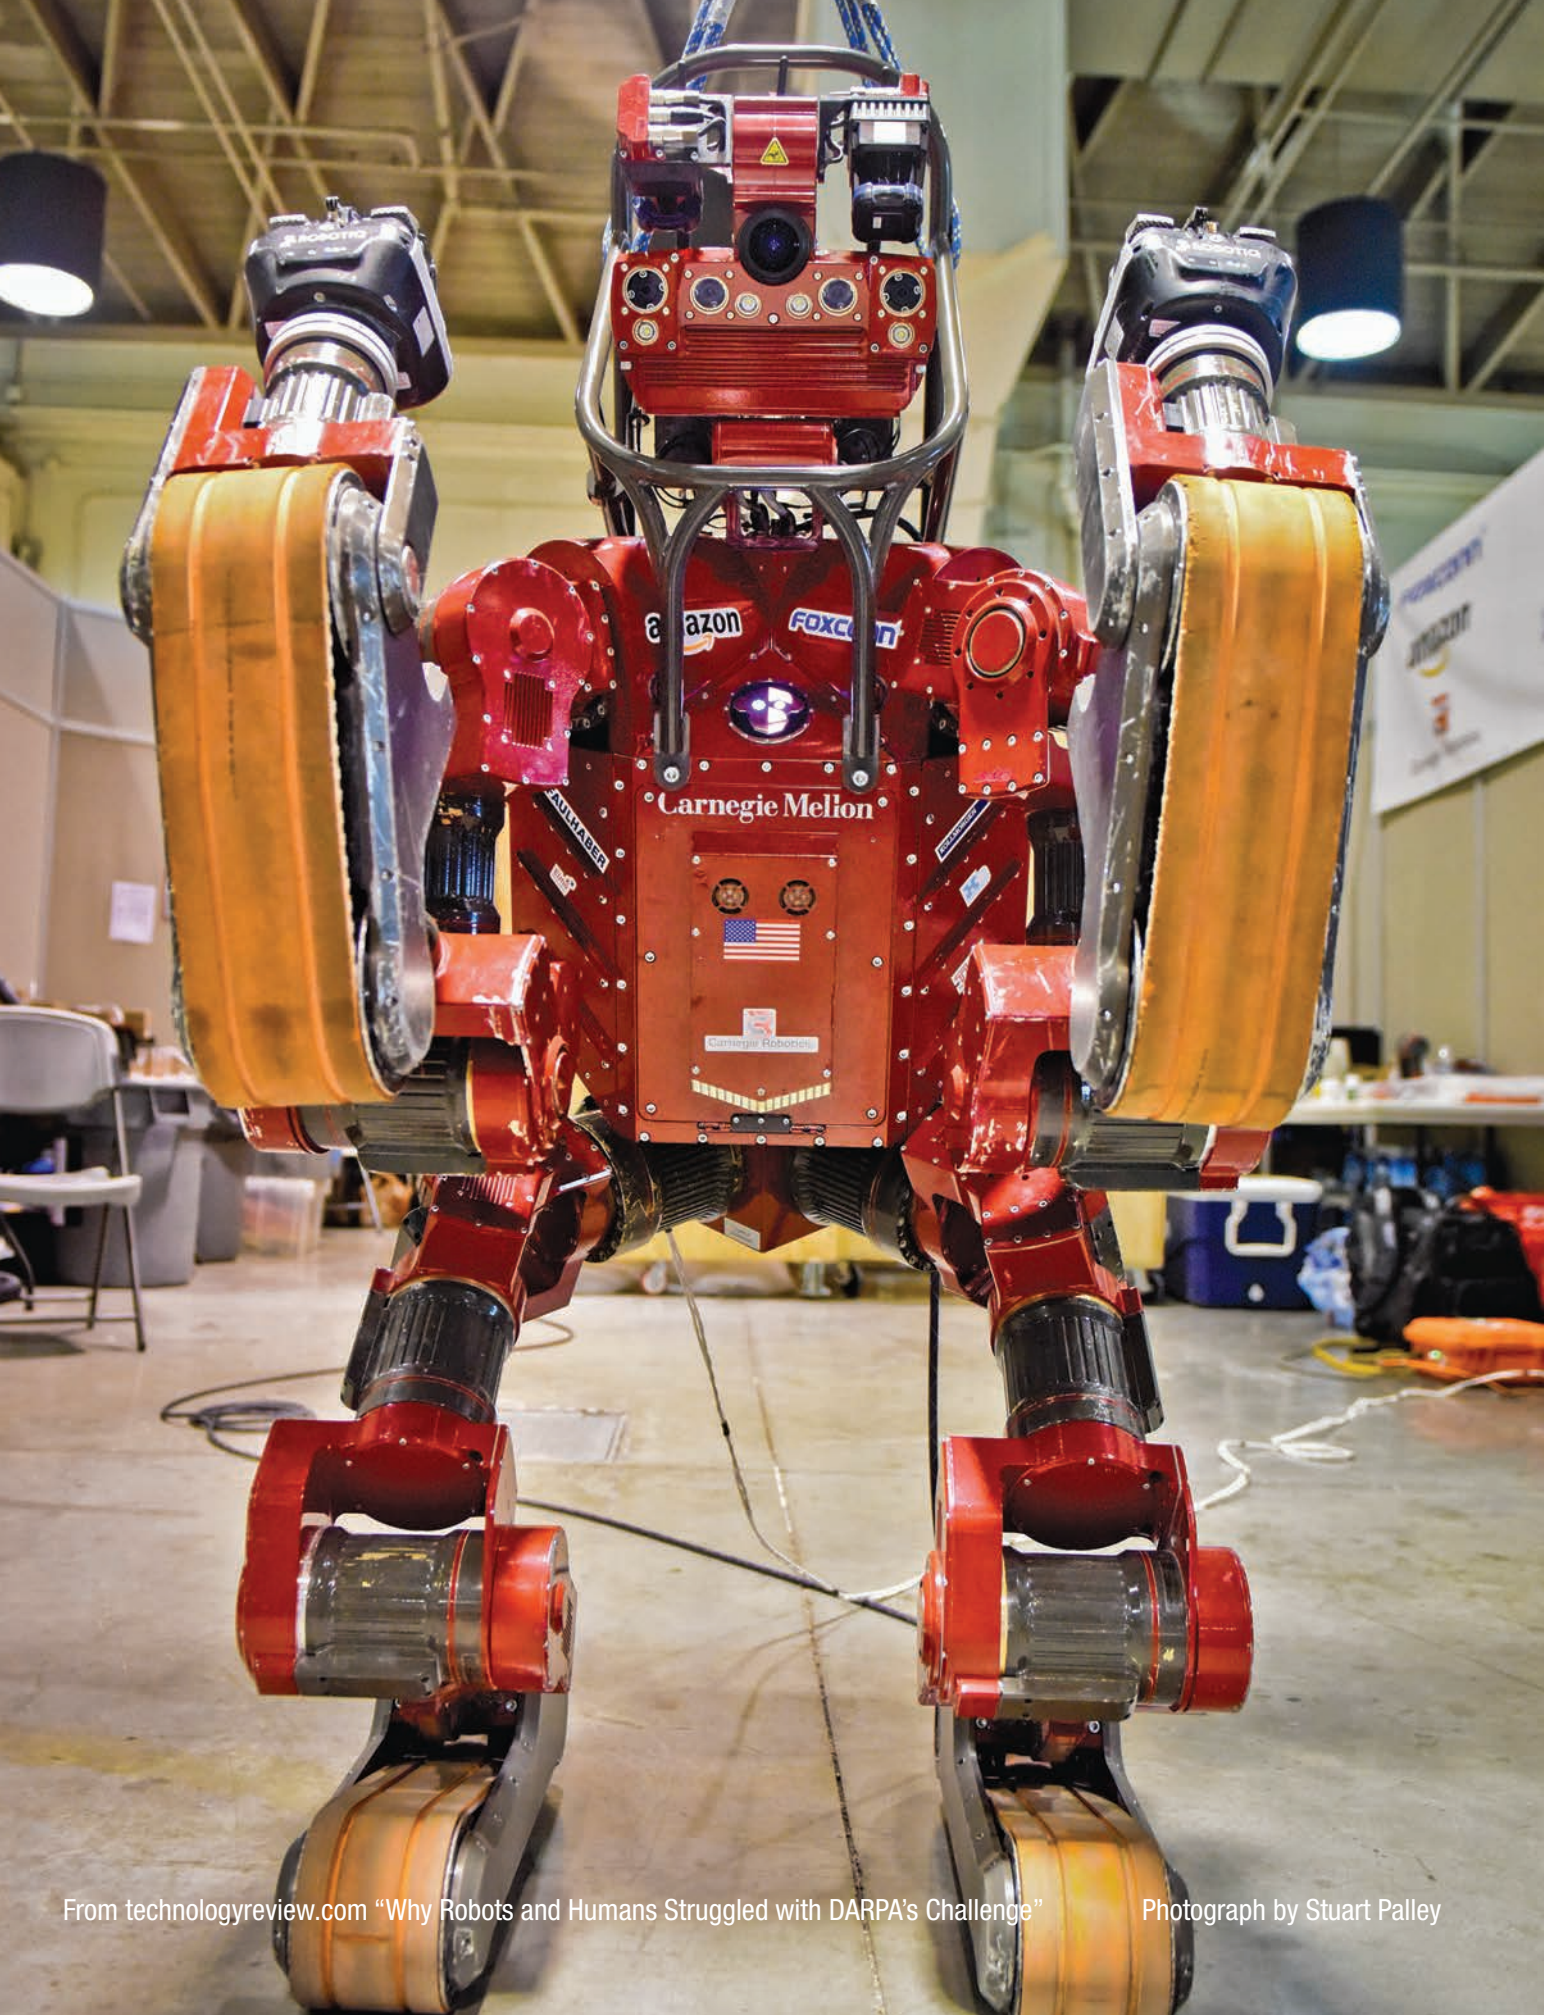

From technologyreview.com “Why Robots and Humans Struggled with DARPA’s Challenge” Photograph by Stuart Palley

Reviews

MIT TECHNOLOGY REVIEW
VOL. 121 | NO. 2

DANIEL ZENDER

Facebook's app for kids should freak parents out

Messenger Kids, its first grab at the under-13 crowd, is not to be trusted. After all, you've seen how the company treats adults.

By Rachel Metz

Do you trust Facebook to take care of your kids?

That's what the world's largest social network is asking parents with the release of its first app for children, Messenger Kids. It's a pint-size version of Facebook's chat app, Messenger (which, like Facebook itself, is intended only for those 13 and older). With Messenger Kids, Facebook becomes the first of the major social networks to put out an app specifically for children under 13.

The move makes sense in some ways. Chat apps are everywhere, so why keep them out of the hands of children? It could even help parents teach them about online etiquette.

And it's a no-brainer for Facebook, whose teenage users are becoming increasingly enamored of competing apps like Snapchat, Twitter, and Kik. With Messenger Kids, perhaps, Facebook can hook younger children on its brand.

The app seems to be part of Facebook's response to the concerns voiced by a growing number of people—including early Facebook investors like Sean Parker, as well as former executives—about the social network's powers of manipulation. In controversial experiments, it has proved it can alter people's moods or their

likelihood of voting by tweaking their news feeds. And during the most recent US presidential election campaign, the company admitted, Russian-backed political content reached 126 million American Facebook users, showing how vulnerable the platform is to abuse.

In an attempt to assuage worries about fake news, Facebook recently announced it will show regular Facebook users fewer posts from companies and news media and more from their friends. With the new app, it's creating a purportedly safe space for children, too. Messenger Kids is meant chiefly for texting and video-chatting with parent-approved friends and family, and it includes fun features like digital stickers and animated masks.

But do I want my young child to use Messenger Kids? The answer is the same as some people's relationship status on Facebook: it's complicated. It's good that children have a chance to grow gradually into social apps, but I don't necessarily want Facebook to be the teacher.

Kids and apps

For a long time, in order to avoid running afoul of a US federal law, the biggest social networks have simply made it against their rules for anyone under 13 to sign up. Called COPPA, or the Children's Online Privacy Protection Rule, the law restricts how companies collect, use, and share online data from kids and aims to give parents control over this data gathering.

But COPPA is clearly ineffective. Three out of five American parents in a 2017 poll conducted on behalf of Facebook and the National PTA (one of the groups Facebook consulted while building Messenger Kids) said that their under-13-year-olds use messaging apps, social media, or both. If you take into account sneakier kids and more oblivious parents, the real figure is likely to be much higher.

The design of Messenger Kids is evidently meant to put parents at ease. It's full of parental controls; kids must get a parent's authorization (via the parent's own Facebook account) to sign up and to add each new contact.

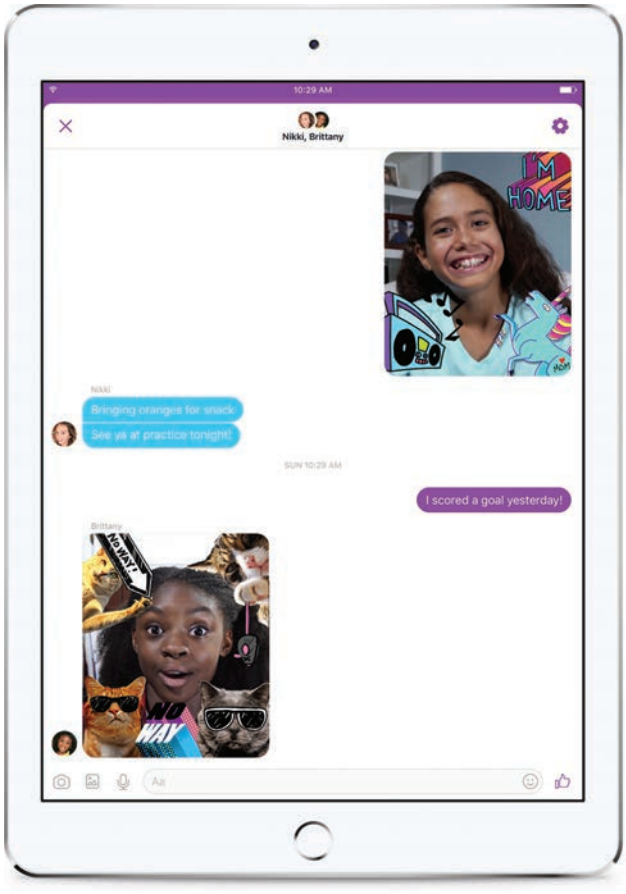

Facebook's Messenger Kids app lets children under 13 do things like text friends and send images with silly frames.

However, the app also has some of the very grown-up features you find on Messenger. For instance, if you send a contact a message on Messenger Kids, it lets you know if the person is online or how long it's been since he or she was active. It will also tell you whether the recipient of your message has viewed it already and if so, for recently sent messages, when.

That kind of information can cause obsessive anxiety even in adults who've already spent years using apps. And habituating kids to always-on communication concerns Tristan Harris, a former design ethicist at Google who cofounded and runs the Center for Humane Technology. "It's like Coca-Cola inventing a kids' soda product," he says. "It still has to sell sugar; it can't really be genuinely concerned with the well-being of kids."

Social cues and tech blues

Researchers at San Diego State University and Florida State University recently found that teens who spent a lot of time using smartphones for things like social media were more likely to be depressed. The work, which shows teen suicide and depression rates climbing in the US since 2010, suggests that these issues are linked to the swift rise in smartphone ownership across the country.

That finding and others like it have alarmed health and education experts. Nineteen groups and nearly 100 individuals (including Harris) signed a letter in January pleading with Facebook to kill Messenger Kids. "Encouraging kids to move their friendships online will interfere with and displace the face-to-face interactions and play that are crucial for building healthy developmental skills, including the ability to read human emotion, delay gratification, and engage with the physical world," the group wrote.

Larry Rosen, a psychology professor and author of *The Distracted Mind: Ancient Brains in a High-Tech World*, says children in the Messenger Kids age group are still honing communication skills, including subtle, nonverbal ones like understanding body language and posture. "I worry we're introducing something to kids who really don't need it," he says, "and I think that the ostensible purpose is wrong, which is to get them started young."

Parent perspective

Not all parents agree. CJ Kanash is an insurance agent in Erie, Pennsylvania, and a father of five. Four of Kanash's children are between the ages of six and 10, and each of them has an Amazon tablet with Messenger Kids on it.

Kanash uses Messenger Kids to send his kids reminders. His nine-year-old son from a previous marriage uses it to say good night to his mom by video chat. "The world is constantly changing, and this is how we're going to communicate more in the future," he says. "They should learn how to do it in a responsible way."

Kids do need to understand how to use technology, and in large part it's up to parents to help them figure out how. And kids under 13 already use social apps, whether we like it or not. But Facebook shouldn't be the default simply because we're familiar with it. The suitable messaging app for young children probably looks nothing like a mini version of Messenger.

For instance, the Children's Design Guide, released in January by a group of designers, researchers, and child and education experts, suggests building digital products that encourage kids to use them in moderation and stop them from sharing potentially harmful information. Meanwhile, Harris and the rest of the letter writers recommend that kids simply use a parent's Facebook or Skype account if they want to connect with far-flung family.

There's little hard research, however, on what a healthy social-media app for kids would look like. As a result, we're running a huge experiment with our children. We haven't had enough time for a generation of kids to grow up with online platforms and report back on how it's affected them. All we have to go on is what we already know: that Facebook's treatment of its older users doesn't bode well for its new focus on kids, no matter how cute and fun the app looks.

Rachel Metz is senior editor for mobile at MIT Technology Review.

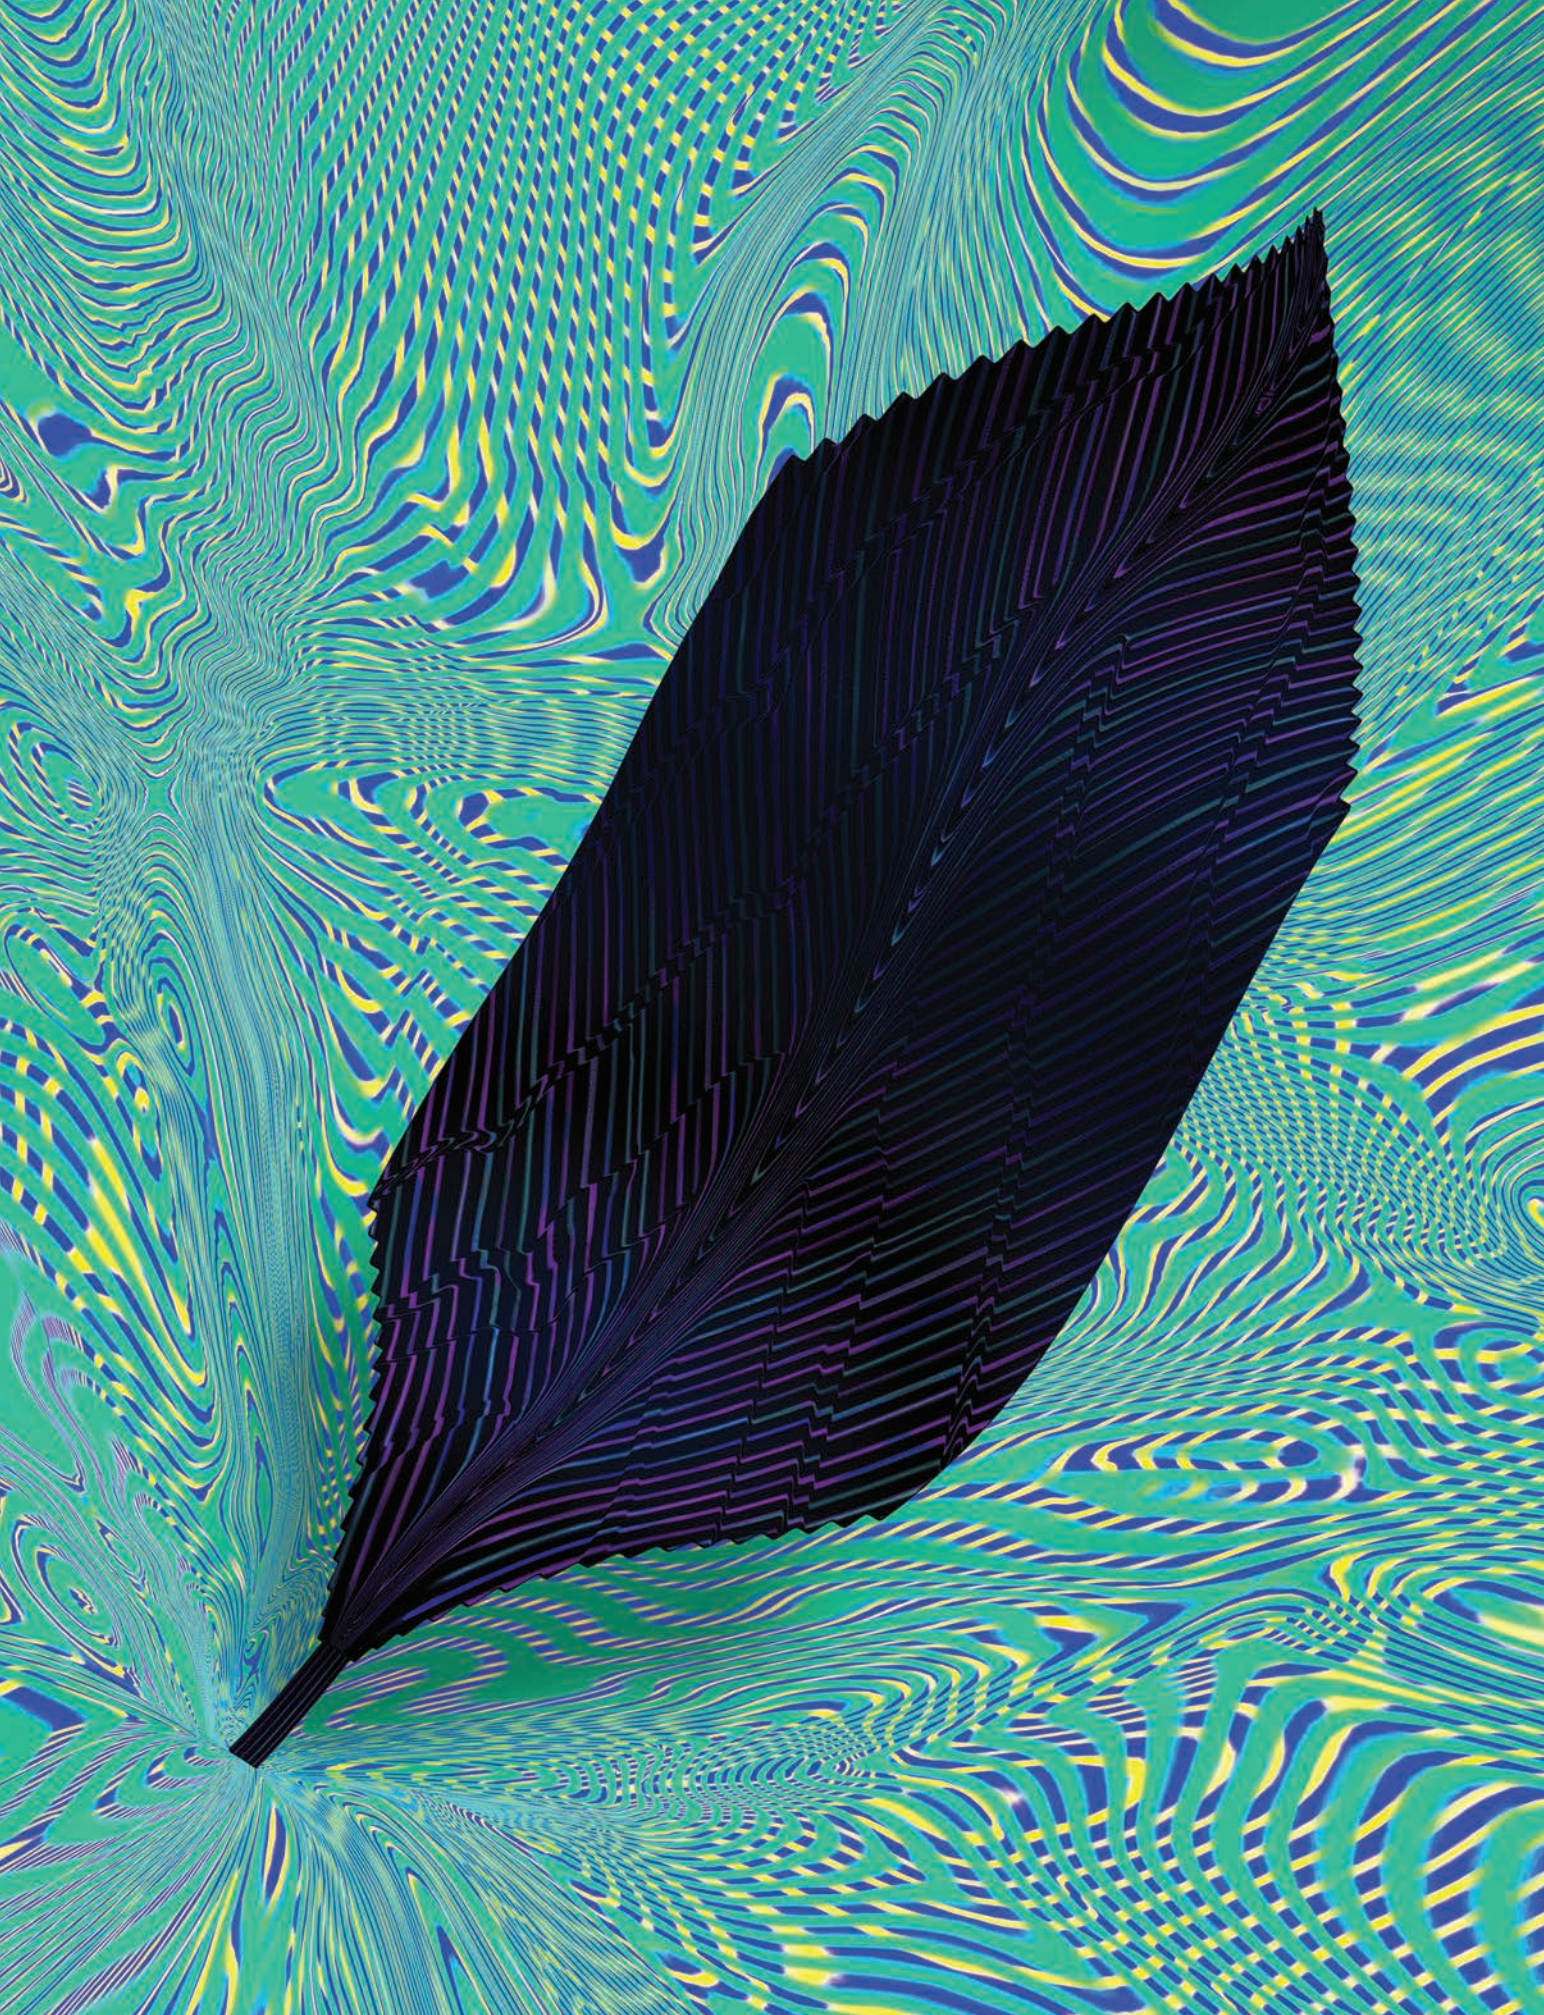

The race to invent the artificial leaf

Since the early 1970s, scientists have been on a quest to develop a technology that could create liquid fuels out of carbon dioxide, water, and sunlight far more efficiently than photosynthesis, the process by which plants harness sunlight to produce carbohydrates and store energy. They call it the artificial leaf.

A commercially viable artificial leaf would solve several of the trickiest challenges in clean energy. It would create a way to directly and affordably store solar energy while producing a carbon-neutral fuel that could transform the transportation sector, even offering a way to make long-distance air travel environmentally sustainable.

Scientists have made slow but considerable progress on the two crucial steps in the process: developing catalysts that use solar energy to split water into oxygen and hydrogen, and creating others that can convert hydrogen and carbon dioxide into an energy-dense fuel. The remaining trick is to combine these tasks in an affordable and scalable way, using cheap and abundant materials.

In this excerpt from a new book, Varun Sivaram, a physicist and fellow at the Council on Foreign Relations, explores the recent progress and diverging research paths of two rival scientists determined to finally deliver and commercialize the artificial leaf: Nathan Lewis at Caltech and Daniel Nocera at Harvard University.

On a balmy Beverly Hills evening recently, members of the Council on Foreign Relations gathered at the Peninsula Hotel to listen to a scientist share his vision for creating an artificial leaf.

Among the collection of executives and ex-ambassadors, most were unsure what to expect. A few exchanged nervous glances as I introduced the speaker's credentials, perhaps bracing for an abstruse physics lecture.

But my guest that night, Nate Lewis, a professor at the California Institute of Technology (Caltech), is a rarity among scientists for his ability to condense complex concepts into memorable soundbites and weave his assorted research strands into a compelling narrative. Lewis, whose gray hair is testament to the decades he's spent chasing an artificial leaf, kicked off his remarks on the future of solar power with this pithy refrain: "Can't store? No power after four."

That failing, he argued in his low, drawling voice, means we urgently need to develop technologies able to store the sun's energy in a fuel that can be used when needed. His preferred route, an integrated solar fuel generator, is an elegant device that takes in water and sunlight and spits out gaseous hydrogen and oxygen. That hydrogen can then be used to fuel vehicles, generate electricity for the grid, or serve as a feedstock to make more complex fuels such as gasoline.

Lewis, also a principal investigator at the federally funded Joint Center for Artificial Photosynthesis, wants his artificial leaf to outperform nature's best plants. Plants, for all their success, are actually terrible at converting sunlight into energy. Even if you don't know anything about how photosynthesis works, you can tell from the leaves' green color that totally efficient energy conversion might not be a plant's top priority (black leaves would be much better at absorbing the sun's rays). The green chloroplasts in leaf cells function well enough for a plant's needs. They perform complex chemical reactions that, fueled by the sun's energy, turn carbon dioxide and water into the energy-storing sugars needed for such activities as surviving and reproducing. When all is said and done, the most efficient plants convert barely 1 percent of the incoming sunlight into stored energy.

Vegetation nonetheless offers a generic model for turning sunlight into fuel. Early on in photosynthesis, plants split water and generate hydrogen and oxygen. The oxygen goes into the atmosphere, while the hydrogen feeds into subsequent chemical reactions.

The way plants accomplish this water-splitting is instructive. The first lesson is that they separate the two halves of the water-splitting chemical reaction—that is, the "half-reactions" that produce hydrogen and oxygen. Evolution was no pyromaniac, and this design choice prevents hydrogen from spontaneously combusting in the presence of oxygen. Second, the plant contains catalysts, or molecules that speed up the half-reactions. Third, plants separate the two half-reactions with a membrane that not only keeps hydrogen and oxygen apart, but also allows charged ions to pass through it, which is important to avoid an imbalance of charge.

Researchers developing solar fuel generators likewise need to put together a similar set of components. Two materials known as "photoelectrodes" are immersed in water and absorb light energy to perform each of the two half-reactions to split water. Two catalysts speed up each of those half-reactions. And a membrane stops the whole contraption—called a "photoelectrochemical cell" (PEC)—from exploding.

But the similarities end there. As Lewis likes to say, after taking inspiration from feathered birds, humans ditched the feathers and invented the 747. Unlike plants, the solar fuel generators

SANTTU MUSTONEN

of the future probably will not use two green photoelectrodes that compete with each other to absorb the same part of the sun's spectrum. Rather, one of them—the anode, which creates oxygen from water—should harness colors of light toward the blue end of the spectrum, and let the colors toward the red end of the spectrum pass through to be absorbed by the cathode below, which produces hydrogen.

Producing affordable energy will require extremely cheap and abundant materials. But that's not all the PEC has to do. To succeed it really needs to be not only cheap, but also safe, robust, and efficient. Unfortunately, so far researchers have only managed to create devices with no more than three of those four characteristics.

Start with safety. To prevent hydrogen and oxygen from combining and exploding, a PEC needs a membrane that separates the two half-reactions. But the half-reaction that produces oxygen from water also turns that water acidic, whereas the half-reaction that produces hydrogen turns nearby water basic. Scientists have to find materials for photoelectrodes and catalysts that do not get dissolved or corroded in acidic or basic media. That demand rules out many cheap materials that wouldn't survive under such conditions. Therefore, making a solar fuel generator out of cheap materials and equipping it with a membrane to ensure safety can lead to it failing the robustness test.

Next, consider the amount of the sun's energy that the device converts into energy stored as hydrogen. That efficiency depends on how well the photoelectrodes collectively absorb sunlight and how fast the two half-reactions split water. With carefully chosen photoelectrodes and catalysts, a solar fuel generator can theoretically achieve more than 30 percent efficiency. Expensive semiconductors offer a diverse buffet of materials to choose from, but cheaper compounds present a far more limited menu. Similarly, precious metal catalysts such as platinum are great at speeding up reactions, but they are rare and costly. The interdisciplinary team of researchers that Lewis led proceeded to throw massive computational power at the problem of finding materials that could satisfy all four criteria, systematically simulating thousands of compounds and testing out the most promising candidates in the lab.

Good old scientific intuition also played an important role in the research process—as has a little luck. Two examples stand out. First, Lewis and his collaborators found inspiration in the catalysts used in oil refineries to strip the air-polluting sulfur out of petroleum products. These catalysts are cheap, and excel at speeding up the half-reaction that produces hydrogen. (Unfortunately, researchers are still looking for a cheap, effective catalyst for the oxygen-producing half-reaction.)

Second, researchers in Lewis' lab accidentally coated their samples with a thin layer of titanium dioxide and found a surprising result. Titanium dioxide is the key ingredient in sunscreen, which protects your skin by blocking ultraviolet rays of sunlight. Here, though, the ultrathin coating played a totally different role, protecting the photoelectrodes and catalysts from being eaten away by the basic solution.

Together, the borrowed insight from the oil industry and the accidental sunscreen discovery allowed Lewis and fellow researchers at Caltech to make a breakthrough. In 2015, they announced an integrated solar fuel generator that was over 10 percent efficient at converting sunlight into hydrogen fuel. The efficiency itself wasn't any great leap—others had reached 22 percent efficiency. But the Caltech device used cheap, Earth-abundant catalysts, and it was able to pump out hydrogen over two days of continuous operation. As a proof of concept, it teased the possibility of a commercially viable product down the road.

If and when this technology leads to a commercial product, it is unlikely to look anything like the leaves that inspired it. Lewis envisions a tarp, rolled out across a vast expanse to soak up the sun's rays, with drainpipes to collect the hydrogen that it produces. That is a far cry from the one-cubic-centimeter prototype his team created, but, listening to Lewis' vision, it's hard not to dream big.

The holy grail

Across the country from Nate Lewis, another acclaimed scientist is also on a quest to commercialize an artificial leaf. Like Lewis, Dan Nocera at Harvard University deftly combines science and communication, moonlighting as a science celebrity—something of a Carl Sagan for solar fuels. He has a knack for connecting with diverse audiences, from American Physical Society scientific gatherings to the Aspen Institute's hobnobbing summits. If his crowd is dining on steak, he'll warm them up by asking: "What did you just chew? The sun! The beef was just the energy of sunlight."

Though both Lewis and Nocera share gray hair, the ability to engage broad audiences, and the same supervisor when they were in graduate school, their approaches to realizing an artificial leaf are dramatically different, resulting in a spirited professional rivalry. Whereas Lewis is laser-focused on producing hydrogen, Nocera wants to leapfrog hydrogen and build a device that harnesses sunlight to directly produce convenient, carbon-containing fuels that can immediately replace today's petroleum products.

For a time, Nocera was content to focus just on hydrogen production. In 2011, he seized the scientific world's attention by plopping what looked like a dark postage stamp into a glass of water, causing hydrogen and oxygen to bubble up on either side. Despite its simplicity, his artificial leaf was the culmination of 30 years of research,

Foundational inventions that change entire industries.

At Qualcomm, inventing comes first. When we connected the phone to the internet, our foundational inventions created the mobile revolution. Now, as we lead the world to 5G, our inventions are going to enable new industries to be created, and the next great product the world can't live without.

qualcomm.com/weinvent

Qualcomm

Inventing the tech the world loves

reaching back to his days as a graduate student at Caltech. Having made the breakthrough, Nocera set out to bring his new technology to market.

Unfortunately, he was about to learn the lesson that almost every other Silicon Valley clean energy startup has learned: the really hard part comes *after* making an exciting lab discovery. He would later lament, "I did a holy grail of science. Great! That doesn't mean I did a holy grail of technology. And that's what scientists and professors don't get." His startup, Sun Catalytix, ended up pivoting away from solar fuel to develop batteries to store energy for the power grid instead (Lockheed bought the company for an undisclosed amount in 2014).

But the experience hasn't stopped him from chasing holy grails, so now, Nocera is pursuing the even harder target of harnessing sunlight, water, and carbon dioxide to produce carbon-based liquid fuels. The prospective benefits of such a technology are compelling. Liquid fuels already have enormous global infrastructure networks, including storage facilities, transcontinental pipelines, and oil supertankers, not to mention ubiquitous filling stations around the world. A device that could transform sunlight into fuels that are already commonly used could piggyback on that infrastructure.

Lewis maintains that the most promising route for making carbon-based fuels from sunlight involves solar-generated hydrogen as an intermediary. From there, well-understood industrial processes could combine hydrogen with carbon dioxide—captured from factories and power plants that burn fossil fuels—to produce a range of useful fuels known as hydrocarbons. A "solar refinery" could create the same range of hydrocarbon fuels that are produced today in oil refineries, and then employed as transport fuels or converted into a range of products from plastics to pharmaceuticals.

As futuristic as all of that sounds, Nocera wants to do something even harder. He wants to bypass the intermediate production of hydrogen and use sunlight, water, and carbon dioxide to *directly* produce carbon-containing fuels. If this maneuver could be done cost-effectively and at scale, it would be the most efficient, single-shot method of storing sunlight in the most versatile fuels known to humankind.

From a scientific point of view, this task looks nearly impossible. Just splitting water to generate hydrogen and oxygen is hard enough. But to create the simplest hydrocarbon—single-carbon methane, which makes up natural gas—is a far more complex proposition. That will require the discovery of even more new materials to absorb light and catalyze chemical reactions. As a result, a commercial technology to make carbon-based fuels directly from solar energy is much further away than one that can produce hydrogen.

Nevertheless, over the last three years, Nocera has made a string of improbable breakthroughs. The first was a conceptual shift: instead of using manmade devices to beat photosynthesis, why not harness nature instead? Nocera knew that nature uses intricate enzymes as catalysts in photosynthesis to convert sunlight into complex sugars. He realized that genetically engineered bacteria could behave similarly after being equipped with an arsenal of potent enzymes.

So in 2015, Nocera built a hybrid device that first split water using an inorganic catalyst to make hydrogen, as other artificial leaf technologies do. The same device then fed the hydrogen, along with pure carbon dioxide, to bacteria, which produced liquid fuels. But although the bugs were terrific at converting carbon dioxide and hydrogen into a variety of fuels, they were incompatible with the inorganic catalyst, which produced forms of reactive oxygen that destroyed the bacteria's DNA.

Then in 2016, Nocera and colleagues published a paper in the journal *Science* triumphantly announcing a new catalyst, made from a cobalt-phosphorus alloy. It not only left the bacteria unharmed, but also self-assembled out of solution, mimicking the self-healing catalysts found in nature. With the catalyst and bacteria working together in harmony, Nocera's device was able to achieve 10 percent efficiency in converting sunlight into alcohol fuels. Nocera reported that the bugs should be able to produce several other carbon-containing molecules for a range of applications from fueling vehicles to producing plastics. And he followed this up by demonstrating in 2017 that a hybrid catalyst-plus-bacteria approach could fix nitrogen in the atmosphere to produce ammonia. That is a tantalizing discovery because over 1 percent of global energy is used today in the production of ammonia to fertilize crops and feed the world. Nocera's prototype suggests that one day, sunlight could power that process rather than fossil fuels.

The jury is still out on whether Nocera's decision to harness living organisms is a good idea. Indeed, bacteria are quite finicky, sensitive to the acidity and temperature of their environment, and thus tough to design around. Smart money, for now, is on devices that harness sunlight to produce hydrogen advancing faster than those that try to directly produce carbon-based fuels. But, by combining modern materials with nature's wizardry, researchers may yet leapfrog simple hydrogen in pursuit of a viable route to the ultimate holy grail: 100 percent clean, drop-in replacements for fossil fuels.

Excerpted from Taming the Sun: Innovations to Harness Solar Energy and Power the Planet *by Varun Sivaram, published by the MIT Press.*

PROFESSIONAL EDUCATION

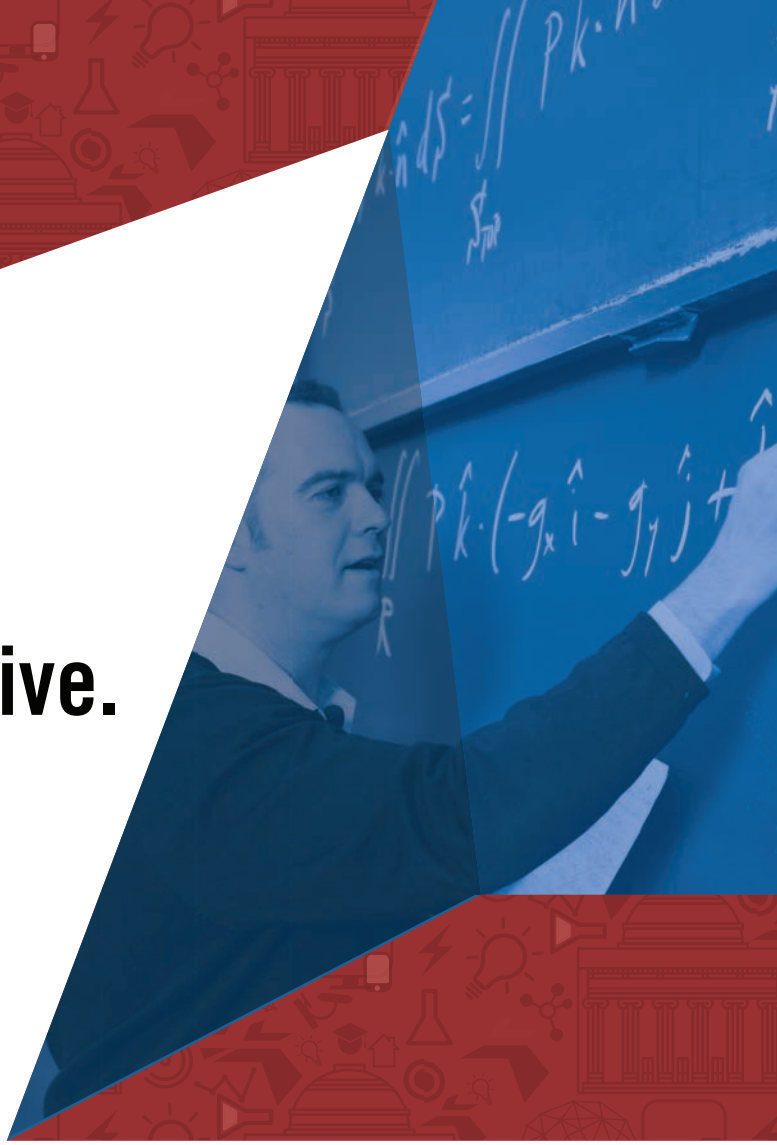

Master. Innovate. Thrive.

Technology professionals from around the world come to MIT Professional Education for the latest research-based knowledge in **biotechnology** and **machine learning** to **manufacturing** and **innovation**. Whether on campus or online, our courses are taught by renowned MIT faculty who are experts in their fields.

UPCOMING COURSES

FEB 12
The Intersection of Leadership & Innovation
Online

FEB 20
Climate Change: Adaption, Mitigation, and Technological Solutions
Online Webinar

MAR 12-13
People Analytics: Transforming Management with Behavioral Data
Japan* *(*in Japanese)*

MAR 28-29
Designing Efficient Deep Learning Systems
Mountain View, CA

PROFESSIONAL CERTIFICATES

We offer professional certificates in the subject areas of **Innovation, Leadership, Machine Learning/Artificial Intelligence** and **Real Estate**.

Learn more and visit: professional.mit.edu/techreview2018
or email: professionaleducation@mit.edu

Demo

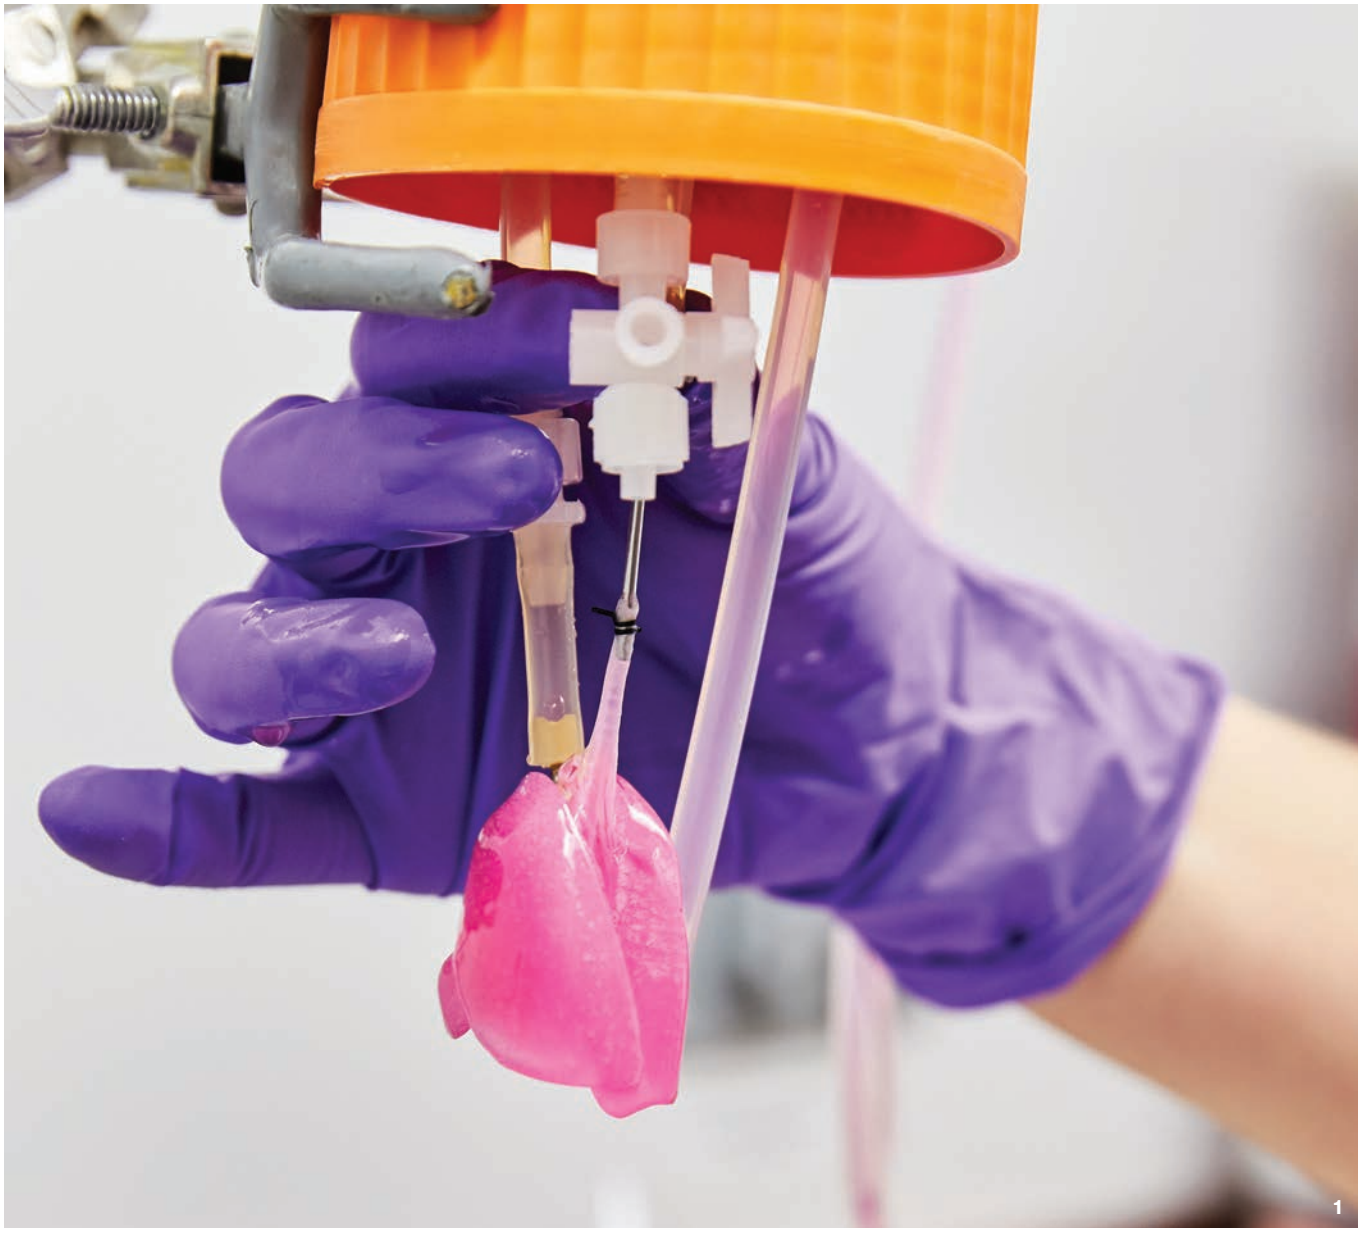

Rebuilding organs so that they might be used again

Harvard's Ott Lab tries to convert animal organs into human-compatible ones.

By Erin Winick
Photographs by Ken Richardson

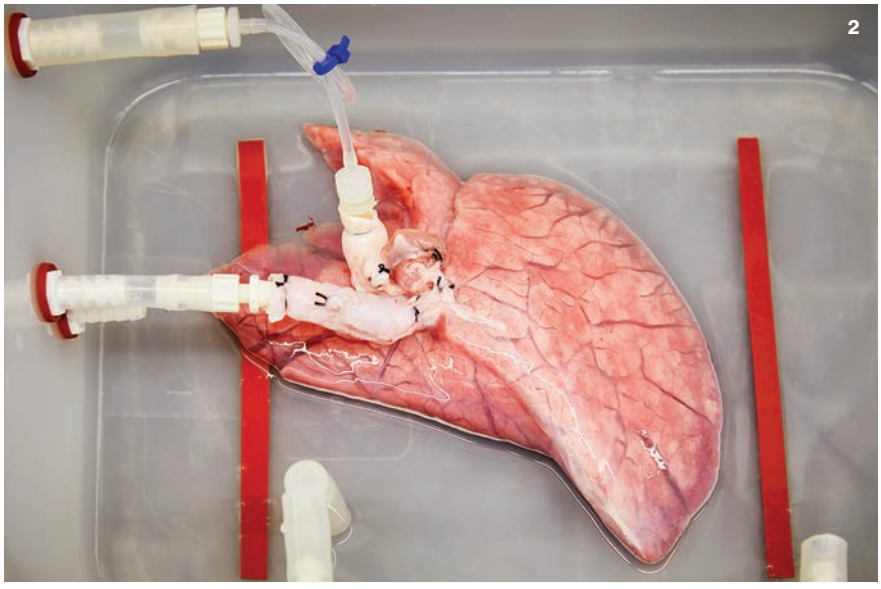

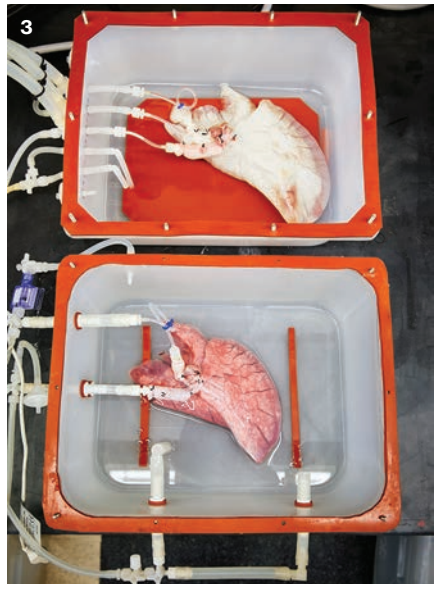

1 Ott Lab researchers pump human cells into a rat lung that has had its own cells removed.

2 The process begins by hooking the circulatory system of an animal organ, in this case a pig lung, to a pump that sends detergent through the organ.

3 As the detergent works its way through, the animal cells are washed away. This changes the color of the lung from pink to white.

4 The lab has also used the same process to successfully remove cells from human lungs, as shown here, for use as scaffolds.

Harald Ott, a surgeon at Harvard Medical School, thinks that his lab's unusual methods might someday solve the organ-transplant crisis. On average, 20 people in the United States die every day awaiting donor organs for transplant, according to the American Transplant Foundation. If Ott's idea works, it could one day eliminate the need for an organ waiting list.

The lab chemically treats organs from rats and pigs to remove all their cells, leaving behind an empty scaffolding. The researchers then populate the scaffold with stem cells from a patient, hoping to create an organ that the human body will accept.

In the most recent work, the lab is using a technique Ott helped develop called lung regeneration, transforming lungs from rats and pigs into human-compatible ones. Though rat lungs aren't suitable as human transplants, they provide an excellent means to test the regeneration procedure.

5 The first organs the Ott Lab used for testing came from rats. Rat lungs like these are still used because their cells can be removed in just two hours.

6 Three rat lungs are tested in an incubator under different conditions to find the best process for regenerating the organs.

Ott originally began to develop this technique while studying at the University of Minnesota (see "Creating a Heart," May/June 2008). In the decade since then, his lab has made strides in deciphering the conditions under which stem cells develop into functioning organs.

So far the team has been able to transplant organs re-created with human stem cells back into pigs and rats. Given that the organs have human cells foreign to the animals, they stay alive for only about a week. But the experiments are evidence that these organs can work in a living organism.

Though the early results are promising, tests in humans are still a ways off. The lab also hasn't yet determined if pig lungs or human lungs will ultimately be the optimal source for scaffolds.

"I feel something we all share here is the belief this really could work."

—Sarah Gilpin, researcher at the Ott Lab

7 A rat lung sits in an incubator where a culture medium is pumped through the organ, growth factors are added, and oxygen is introduced to make the lung "breathe."

In a first, humans launch an electric car into space

When SpaceX launched its Falcon Heavy rocket in February, a lot of attention was devoted to its payload—Elon Musk's personal cherry-red Tesla Roadster, "driven" by a spacesuit-wearing mannequin. The roadster and mannequin will now orbit the sun together, but not forever—just for the next few hundred million years, Musk estimates.

SPACEX

MIT Technology Review (ISSN 1099-274X), March/April 2018 issue, Reg. US Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2018. The editors seek diverse views, and authors' opinions do not represent the official policies of their institutions or those of MIT. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: Send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the internet at www.technologyreview.com/customerservice. Basic subscription rates: $41.94 per year within the United States; in all other countries, US$54. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051, Fort Erie, ON L2A 6C7. Printed in USA. Audited by the Alliance for Audited Media

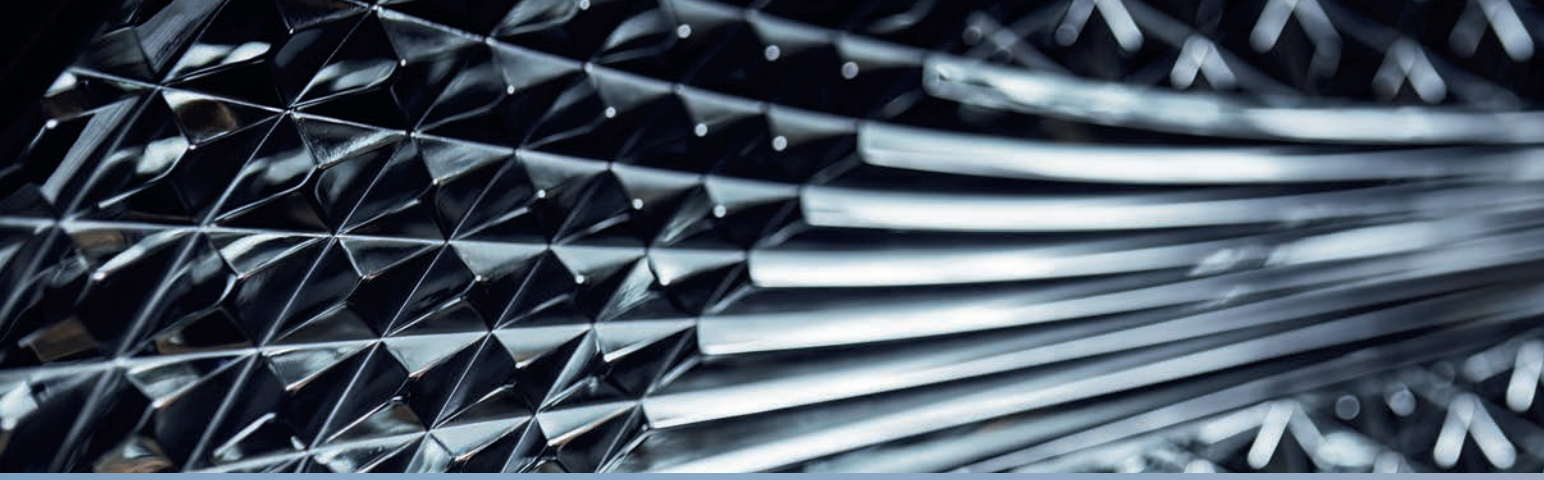

RARE FORM

THE FIRST-EVER LS 500. LIVE IN THE NEW.

We didn't merely try something new. We crafted an entirely new experience. With accents like *Kiriko* glass, designed by hand then delicately etched with thousands of cuts to catch the eye and transform in the light. A 416-horsepower[1] twin-turbo engine paired with the first-ever 10-speed Direct-Shift automatic transmission in its class[2] takes you from 0 to 60 in just 4.6 seconds.[1,3] All this is complemented by cutting-edge technology with one of the largest Heads-Up Displays in the industry.[2] The LS 500 isn't simply new. It's the most evocative LS ever.

LEXUS
EXPERIENCE AMAZING

lexus.com/LS | #LexusLS

Options shown. 1. Ratings achieved using the required premium unleaded gasoline with an octane rating of 91 or higher. If premium fuel is not used, performance will decrease. 2. 2018 LS vs. 2017/2018 competitors. Information from manufacturers' websites as of 8/4/2017. 3. Performance figures are for comparison only and were obtained with prototype vehicles by professional drivers using special safety equipment and procedures. Do not attempt.

This device has not been authorized as required by the rules of the Federal Communications Commission.
This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.